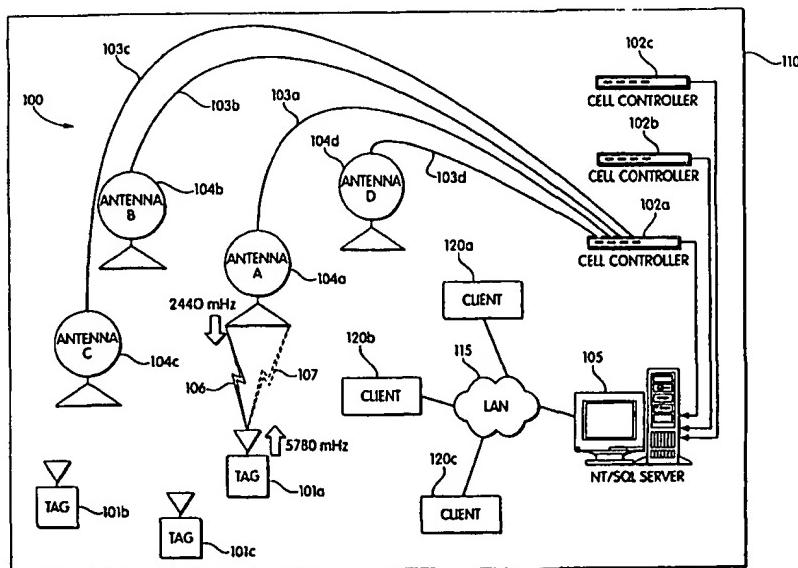




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G08B 5/22, H04Q 5/22, H03H 7/30, G01S 13/74	A1	(11) International Publication Number: WO 00/46771 (43) International Publication Date: 10 August 2000 (10.08.00)
(21) International Application Number: PCT/US00/02633		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 1 February 2000 (01.02.00)		
(30) Priority Data: 09/244,600 4 February 1999 (04.02.99) US		
(71) Applicant: PINPOINT CORPORATION [US/US]; One Fortune Drive, Billerica, MA 01821 (US).		
(72) Inventors: BRADY, David; 142 Hinckley Road, Milton, MA 02186 (US). LANLZ, Colin; 83 Nottingham Drive, Nashua, NH 03062 (US). LOOMIS, Gary; 230 Prospect Street, Leominster, MA 01453 (US). WERB, Jay; 44 Lombard Street, Newton, MA 02458 (US).		Published <i>With international search report.</i>
(74) Agents: PRITZKER, Randy, J. et al.; Wolf, Greenfield & Sacks, P.C., Federal Reserve Plaza, 600 Atlantic Avenue, Boston, MA 02210 (US).		

(54) Title: TECHNIQUE FOR FILTERING SIGNALS IN A LOCAL POSITIONING SYSTEM



(57) Abstract

System (100) for tracking mobile tags (101a-c). Cell controllers (102a-c) with multiple antenna modules (104a-d) generate a carrier signal (106) which is received by the tags. Tags shift the frequency of the carrier signal, modulate an identification code onto it, and transmit the resulting tag signal at randomized intervals. The antennas receive and process the response, and determine the presence of the tags by proximity and triangulation. The recursive-least squares (RLS) technique is used in filtering received signals. Distance of a tag from an antenna is calculated by measuring the round trip signal time. The cell controllers send data from the antenna to a host computer (105). The host computer collects the data and resolves them into positional estimates. Data are archived in a data warehouse, such as an SQL server.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

TECHNIQUE FOR FILTERING SIGNALS IN A LOCAL POSITIONING SYSTEM

Background of the Invention

The invention relates to tracking systems, and more particularly to tracking systems designed to track articles and personnel.

Various techniques are used to locate and track the positioning of, for example, personnel and objects. Global positioning system (GPS) technology is generally used to locate objects over a large outdoor area by using received signals generated by satellites. Techniques included in the GPS technology are not generally readily applicable for use in tracking objects locally over a smaller area. One problem with using GPS techniques for local tracking relates to the accuracy of locating an object. For GPS applications, accuracy in the range of 50 meters is ample. For tracking objects locally, greater accuracy, for example accurate location of an object within 10 meters, is generally needed. Additionally, local tracking applications may include tracking objects and people indoors, whereas GPS systems are generally employed in outdoor applications. In indoor applications, indoor multipath effects may be encountered. When trying to locate an object indoors, the indoor multipath effects introduce additional interference signals which must be filtered from a return signal to accurately locate an object. Indoor multipath effects may be generally characterized as the "noise" or interference introduced in a received signal from being indoors. Generally, "multipath" refers to the phenomenon exhibited by a transmitted signal when it takes multiple routes between a transmitter and a receiver, such as a transmitted signal which may be reflected from a metal beam in a building. Additional types of "indoor noise", such as the tones generated by a microwave oven, are also generally filtered or separated out of a received signal to be able to accurately locate an object. The GPS techniques generally do not take into account these indoor multipath effects or other types of indoor noise. Rather, the GPS techniques may take into account different types of interference signals and other effects which are generally not a concern for local position systems.

Other techniques exist which employ a variety of short-range radio-based technologies used to track items indoors. These systems generally identify objects with a sensor having a range of a few centimeters to about 3 meters, depending upon the technology. For example, electronic article surveillance (EAS) systems are widely used in the retail industry to locate items. Generally, in EAS systems the tags respond to an

electronic field by resonating and when the resonance is detected, an alarm sounds. However, EAS products generally have a restricted range.

Another newer technology, radio frequency identification (RFID), has increased in use in recent years. RFID tags are identified as they pass fixed sensors. Several RFID systems use backscatter to communicate by which a tag responds to an interrogator signal by modulating the response signal and re-radiating or transmitting it at the same carrier frequency. However, these systems generally have a limited range of about 3 meters, and additionally, the backscattering technique used in tag communications may be unreliable at ranges greater than 3 meters due to the signal reflections from, for example, filing cabinets, white boards, and the like.

Another current tagging system is the Electronic Home Arrest Monitoring system used to monitor persons under house arrest. In this application, a radio-frequency transmitter around a person's ankle emits a signal which is received by a monitoring device. The signal includes an identification number and other information.

Yet another type of system employs infrared technology, called IRID, for locating and identifying objects. Tags in these types of system periodically transmit their identification codes by emitting infrared light to a receiver or reader. IRID systems do not work well in all types of local positioning applications. For example, IRID systems may not function properly under various common lighting conditions. Additionally, for example, if a tag is used to locate people, a user's cooperation is generally needed to avoid blocking a tag's IR emitter.

Techniques employed in existing systems generally do not efficiently locate objects to a great enough accuracy while simultaneously taking into account the various types of interferences that may be encountered in a wide variety of local positioning applications.

Thus, there is required an efficient technique for locating, with a high degree of accuracy, people and objects in local positioning applications.

Summary of the Invention

In accordance with principles of the invention, disclosed is a method of filtering received signals in an article tracking system to determine a received tag signal. The distance of an object may be determined using the difference between the time the tag signal is received and the time the tag signal is transmitted. However, in various environments, there is often a problem in determining which received signal is the received tag signal due

to various noise components which confuse identification of a received signal as the received tag signal. In the method disclosed, a plurality of signals are received, and a first range of these received signals is determined. An autocorrelation function with phase shifts corresponding to the first range of received signals is performed producing a corresponding magnitude for each received signal included in the first range. A second range of the received signals is determined. The second range is a subset of the first range, and the second range approximates where in the first range the received tag signal is located. A recursive least-squares technique is performed using the second range of received signals. This technique produces a vector of filtered signals corresponding to the second range of received signals. A peak corresponding to a filtered portion of the received tag signal is approximated using the vector of filtered signals. The location of a tag associated with the received tag signal is determined using the time at which the received tag signal is sampled.

Thus, there is described an efficient technique for locating, with a high degree of accuracy, people and objects in local positioning applications.

15

Brief Description of the Drawings

Figure 1 shows an overview of a system configured according to the invention; Figure 2 shows several cell controllers deployed in a multi-story building; Figure 3 is a block diagram of a tag RF design according to the invention; Figure 4 is a block diagram of an alternative embodiment of a tag; Figures 5A-5G are diagrams of a signal as it passes through various stages of the system; Figure 6 is a block diagram of the cell controller RF design; Figure 7 is a block diagram of a cell controller active antenna module; Figure 8 is a block diagram of a modulator RF design; Figure 9 is a block diagram of a cell controller cable extender module; Figure 10 is a block diagram of a cell controller; Figure 11 illustrates extraction of tag data from a series of correlations; Figures 12A-C are diagrams of tag datagrams; Figure 13 shows a tag incorporating a delay element; Figure 14 shows several cell controller receive chains operating in parallel; Figure 15A is a block diagram of an embodiment of the signal processing hardware of Figure 10;

Figure 15B is a flowchart depicting method steps of one embodiment of the Signal processing hardware unit of Figure 10;

Figure 16 is a block diagram of an embodiment of a signal filtering technique;

Figure 17 is a block diagram of an embodiment of the transversal filter of Figure 16;

5 Figure 18 is a flowchart depicting method steps of one embodiment of the recursive-least squares (RLS) technique as used in a method step of Figure 15B;

Figures 19A-19E are diagrams of sample waveforms in embodiments of the system of Figure 1; and

10 Figure 20 is a flowchart depicting method steps of one embodiment of approximating a peak of the filtered tag signal.

Detailed Description

The local positioning system described in paragraphs that follow is an article tracking system that may be used, for example, to track and locate objects indoors. This 15 system may be characterized as a radio-frequency identification system. Generally, such a system may also be used in a variety of other local positioning applications, such as outdoor tracking of objects or locating personnel indoors as well as outdoors. Limitations specific to the embodiments described herein are not meant to imply limitations to the claimed invention.

20 Radio Frequency Identification (RFID) products typically have three components: (1) a tag (the item being identified), (2) an interrogator (a device which detects the presence of a tag), and (3) a system (typically including cabling, computers, and software which tie together the tags and interrogators into a useful solution). RFID products are typically designed to detect tags when they pass within range of a few fixed or handheld 25 interrogators.

RFID systems are usually deployed as high-end replacement technology for bar coding. RFID and related systems include passive RFID systems, active RFID systems, infrared ID systems, and Electronic Article Surveillance (EAS) systems.

The tags in a passive RFID system do not carry on-board power. The interrogator in 30 such systems transmits operating power for the tags. Such systems generally have a detection range of a meter or less, although somewhat longer ranges have been achieved. Typically, these systems operate in the 125-kilohertz or 13.56 megahertz radio band.

Most passive RFID systems work as follows. An interrogator emits an electromagnetic field for the purpose of powering the tag. A coil in the tag is powered by the electromagnetic field, causing the tag's circuitry to "wake up." The tag uses this power to send an identifying signal back to the interrogator.

5 Although some passive RFID systems are read-only (that is, the tags in such system respond to a query by reading information from their memory and sending the information back to the interrogator), the tags used in other passive RFID systems have a limited ability to accept information and instructions from the interrogator, for example read/write capabilities in smart cards (electronic money) and "electronic manifests" in industrial applications.

10 Passive RFID tags have been employed in conjunction with access control, smart cards, vehicle identification (AVI), waste management, item tracking, animal identification, manufacturing control, materials handling, and a variety of other purposes.

15 One fundamental design goal of any RFID system is for the weak signal emitted from the tag to be distinguishable from the much stronger signal emitted by the interrogator. Some strategies for doing this include:

Frequency shifting. Circuitry in the tag receives a carrier from the interrogator, translates the signal to another frequency, and emits a response modulated onto that second frequency.

20 Half duplex operation. The tag is charged by the interrogator. When the interrogator's charging circuit turns off, the tag uses the stored power to respond.

Modulated backscatter. The tag modulates its antenna cross section to identify itself to the interrogator.

25 Delayed retransmission. Surface acoustic wave (SAW) devices retransmit the interrogator's carrier after a delay. The tag's identity is indicated by time variations in the delayed response.

Active RFID systems require battery-powered tags. The battery permits a longer detection range of between 3 and 100 meters. These systems are capable of locating tags with higher accuracy than passive RFID systems and typically operate in the 400, 900, 30 2440, or 5800 megahertz bands. Active tags tend to enable multiple tags to be within range of an interrogator by the use of "handshaking" between the tags and interrogator, so that each tag transmits its signal in turn. Communication between tag and interrogator in active RFID systems is also typically faster than with passive tags.

Most active RFID tags respond to the interrogator when polled, in accordance with a communications protocol. Some active RFID tags "chirp" (transmit) a signal spontaneously at predetermined intervals. A tag's chirped signal is detected by the interrogator if the tag is in range of the interrogator.

5 Infrared systems (IRID), while not RFID systems, also endeavor to detect and identify the position of mobile tags. A typical IRID system includes a tag that chirps its identity at randomized intervals. Infrared readers located in the ceiling detect these transmissions, and report them to a host. The transmission rate from the tag to reader is typically about 600 baud. Motion detectors in the tags enable the tags to transmit more
10 frequently when in motion. The tags are typically about the size of dominos.

EAS systems are intended to deter theft in retail environments. EAS tags are fairly unreliable, very low in cost, and limited in capabilities. Although they track mobile tags, they are not generally considered to be identification products, because EAS tags are uncoded and cannot be distinguished from one another.

15 Referring to Figure 1, an article tracking system 100 contains the following general components:

Tags: Inexpensive miniature radio frequency transponding tags 101a-c are attached to people and/or items being tracked. Tags 101a-c "wake up" periodically, and "chirp" (transmit) a radio-coded unique identity code (UID). The tags 101a-c are designed so that
20 their range is 15-30 meters in a typical indoor environment, the range mostly being limited by a need to conserve the life of the tag battery, and the requirement that the tag 101a and tag battery be small and thin.

Cell Controllers: Cell controllers 102a-c detect the chirps of tags 101a-c and calculate the distances of those tags 101a-c to active antenna modules 104a-d connected to the cell controllers 102a-c. Each antenna module preferably has a transmit antenna and a receive antenna. In Figure 1, the antenna modules connected to cell controllers 102b and 102c are omitted for simplicity. A cell controller 102a is typically contained in a case and is mounted behind a hung ceiling. The cell controller 102a can receive power from a conventional wall outlet or the equivalent. The cell controller 102a is attached through
25 coaxial cables 103a-d to the antenna modules 104a-d, respectively, which provide coverage of an area of the indoor facility 110. A tag signal 107, transmitted by a tag 101a, is received by one or more antenna modules 104a-d, and is processed by chips in the cell controller 102a, such as digital signal processing (DSP) chips as may be included in the signal
30

processing hardware. The information resulting from this processing is used to identify both the identity of the transmitting tag 101a and the distance between the tag 101a and each, for example, of the receiving antenna modules 104a-d.

5 Host Computer: Cell controllers 102a-c are in data communication with a host computer 105, which collects data and information available from the cell controllers 102a-c and archives the data into an open format database, such as an SQL Server.

User Application: In a preferred option, client workstations 120a-c communicate with the host computer 105 over a network, such as a LAN 115. A client application running on each of the client workstations 120a-c can access the SQL Server and presents 10 the data in a way that is useful to an end-user.

The tag 101a does not generate its own radio signal. Rather, an antenna module, for example antenna module 104a, continuously transmits a direct sequence spread spectrum interrogator signal 106 at a first frequency, for example 2440-megahertz. The tag 101a receives this signal 106, modulates its UID code onto the signal 106, and immediately 15 transmits back a frequency-shifted signal 107 at, for example, 5780-megahertz to, for example, antenna 104a. The distance from the antenna module 104a to the tag 101a can then be determined by the cell controller 102a from the round trip time of transmission, accounting for fixed and known delays in the wiring and electronics. The cell controller 102a can quickly switch among antenna modules 104a-d to obtain the distance from the tag 20 101a to each of antenna modules 104a-d (which receive the return signal 107), and from that information determines the tag's location by triangulation techniques.

The system 100 is designed to be scaleable, allowing addition of cell controllers to existing cell controllers 102a-c and antenna modules to existing antennas modules 104a-d to improve the precision with which tag location is determined. Figure 2 shows how a 25 collection of cell controllers 102a-c can be deployed in the large multistory building 110. As shown in Figure 2, multiple cell controllers 102a-c feed data to a single host computer 105, typically through a TCP/IP communications network. A variety of data protocols and transfer mechanisms can be used in preferred embodiments. For example, if a local area network is not available, connection to the host can be accomplished via RS485, RS232, 30 RS422, power line modem, or a dedicated phone line. Alternatively, specialized modems designed for use on such cables can be employed.

Each of the cell controllers 102a-c can be installed to cover a separate floor 130a-c, respectively, with the exact configuration being modifiable by a system administrator. On

floor 130a, cell controller 102a, with its collection of antenna modules 104a-d, is installed in the ceiling 140a. The same configuration of equipment is used on the remaining floors 130b-c. Antenna modules 104a-d are designed to provide good gain downwardly and horizontally, and poor gain upwardly, so that a vertical location (that is, the floor) of a tag 5 101a can be determined by noting which antenna modules 104a-d receive the strongest signals from the tag 101a. Structurally, a ground plane is placed behind each antenna to reflect signals downward. The horizontal location of a tag 101a is then roughly determined by noting which antenna modules 104a-d receive a strong signal from the tag 101a. The horizontal location of the tag 101a with respect to an antenna module 104a can be 10 determined more precisely by estimating the distance from the tag 101a to each antenna module 104a-d, based on the combined time of flight of the interrogation signal 106 and the tag signal 107. Each "cell," consisting of a cell controller 102a and its antenna modules 104a-d, covers several thousand square feet of floor space. Each cell operates independently, enabling more cells to be added without affecting the performance of 15 existing cells.

If a user desires to locate tags by "zone," one antenna per zone can be installed. Users wishing to track one or more tags 101a-c moving down hallways, can install antenna modules 104a-d every 20 or so meters along a hallway 130a-c, and calculate the linear location of a tag 101a by measuring the distance from the tag 101a to those antenna 20 modules 104a-d. Customers wishing to triangulate upon the location of a tag 101a must install enough antenna modules such that the tag 101a will be in range of at least three of the antenna modules. A typical installation will cover a complete facility 110 with a combination of "zone" and "hallway" coverage at a relatively low cost per square foot, and, over time, upgrade certain areas with enough antenna modules to triangulate upon a tag's 25 position.

Other preferred embodiments may include a wall mounted antenna radiating generally horizontally, rather than vertically. The selection of antenna type may be based on a variety of functional factors familiar to one of ordinary skill in the art.

30 Tag RF Design

Referring to Figure 3, a tag RF circuitry 300 receives signal 106 at a tag receive antenna 301 and transmits tag signal 107 at a tag transmit antenna 312. The function of the tag RF circuitry 300 is to transpond the incoming spread spectrum signal 106 by frequency

translation. The secondary function of the tag RF circuitry 300 is to modulate tag data onto the emitted tag signal 107, under the control of a microprocessor 308. The information emitted on the tag signal 107 includes, in a preferred embodiment of the invention, the serial number of the tag, datagram headers, and tag data 309 such as that derived from a motion indicator or a low power indicator.

The incoming signal 106 is preferably a direct sequence spread spectrum signal, biphase or quadrature modulated from the cell controller 102a, in the 2440 megahertz band. The signal 106 is received by the tag receive antenna 301, which collects the signal 106 and feeds it into the tag RF circuitry 300.

After the signal 106 is received by the tag receive antenna 301, an Rx (receive) bandpass filter 302 ensures that the tag is receiving signals only in the 2440 megahertz ISM band, rejecting radar signals, electronic newsgathering signals, and the like. In one embodiment, the filter 302 is implemented as an etched coupled stripline filter embedded in the circuit board. The signal 106 is then amplified by an amplifier 303 to ensure that the received signal can be mixed, in a frequency mixer 304, without degrading the signal to noise ratio (SNR).

The frequency mixer 304 translates or shifts the carrier frequency from 2440 megahertz to 5780 megahertz. The incoming signal, with a center frequency of 2440 megahertz, is mixed with the output of a phase locked oscillator (PLO) 305 with a center frequency of 3340 megahertz. This results in a sum frequency of 5780, along with a difference frequency and a variety of harmonics and subharmonics which are removed with a bandpass filter 306. In one embodiment, the PLO 305 consists of a phase locked loop (PLL) chip with three inputs: (1) a sampled output from a voltage controlled oscillator (VCO); (2) a reference tone from a 10 megahertz oscillator; and (3) a frequency programming interface to a microprocessor 308. This generates a pure tone with good phase noise at the 3340-megahertz tag LO frequency. In an alternative embodiment, the PLO 305 outputs a 1670-megahertz tone, which is then doubled to give the desired 3340-megahertz result.

The next element of the tag RF circuitry 300 is a biphase modulator 307 which, under control of the microprocessor 308, can either pass the 5780-megahertz signal unaltered, or change the phase of the signal by 180 degrees. The modulator 307 is implemented as a single pole double throw RF switch 801 that feeds a 180 degree hybrid, as shown in Figure 8. Several forms of modulation can be used, including on-off keyed (OOK)

modulation, binary phase-shift keyed (BPSK) modulation, multiple phase-shift keyed (MPK) modulation, and quadrature amplified (QAM) modulation. BPSK is the preferred form of modulation. The output from the modulator 307 is fed into an amplifier 310, then is filtered by a transmitter bandpass filter 311, and the output of filter 311 is emitted from a
5 transmit antenna 312 as the tag signal 107. Since the amplifier 310 operates at high frequency, it consumes significant power, and alternative embodiments (such as that shown in Figure 4) that make this amplifier 310 unnecessary are preferred. The Tx Filter 311, implemented as a 5-pole filter, is necessary to ensure tag compliance with FCC Part 15 requirements.

10 The tag RF circuitry 300 shown in Figure 3 is intended to illustrate the general functions of a tag 101a-c, with an embodiment that is workable and self-explanatory. Those skilled in the art will be able to combine multiple functions into single elements in order to conserve power and take full advantage of available parts, or implement the same functions with a custom ASIC. Figure 4 shows an alternative embodiment 400 which fulfills the
15 same basic functions as that shown in Figure 3, but with fewer components and using less power. The essential difference between the circuitry 400 shown in Figure 4 and the circuitry 300 shown in Figure 3 is that the modulator 404 in Figure 4 is placed before the frequency mixer 406 in order to reduce the number of components (for example, the amplifier 310 is eliminated) and to conserve power.

20 In the place of the frequency mixer 304 (Figure 3) or the time delay element 1505 (Figure 13), other signal transmission discriminators may be used to transpond by other methods. For example, a tag, such as the tag 101a, can transpond using backscatter, frequency translation by mixing, frequency translation by taking a harmonic, frequency translation by taking a subharmonic, or by signal delay (such as via a SAW device).

25 Not shown in Figure 4, but desirable for the tag RF circuitry, is the use of a common crystal reference for both the PLO 407 and clock timing in the microprocessor 405. Accurate timing is an important, if not critical, feature of the system, enabling the cell controllers 102a-c to predict when a tag 101a will transmit a tag signal 107. Using the same crystal reference in the PLO 407 and in the microprocessor 405 clock timing allows the cell
30 controller 102a to accurately calibrate the source by measuring phase shifts in the received signal (as described hereinafter), and to synchronize its clock timing accordingly.

Not shown in Figure 4, but desirable for some applications, is an embodiment where the transmit antenna 409 and receive antenna 401 are combined into a single element, and which uses a diplexer with the single antenna structure.

The manner in which the tags 101a-c are powered depends on the application. (Note that Figure 3 and 4 omit the tag power source.) Typically, a tag 101a will be battery powered, with the RF stage turned on and off under control of the microprocessor 405. In a preferred embodiment, the microprocessor 405 goes into a low power state, where it merely waits until it is time to power up the tag 101a again. In an alternative embodiment, all tag circuitry 400 cycles on and off under analog control, using an RC time constant in the circuit 400 as the timing source.

Using the tag RF circuitry 300 or 400 of Figure 3 or 4, if a tag 101a is in range of two of the cell controllers 102a-c, and those cell controllers are sending pseudonoise with low cross correlation characteristics, the tag 101a will correctly transpond both signals simultaneously.

Tags 101a-c require a period of time, on the order of a millisecond, to charge up and charge down. During these brief periods, typically, the tags 101a-c will not be stable enough to use, but will nonetheless emit RF into the radio channel through the transmit antenna 409. For high-performance applications, where radio bandwidth is limited, a microprocessor-controlled switch can be added to the tag's transmit chain to eliminate such spurious emissions.

The tag RF circuitry 300, 400 shown in Figure 3 and 4 can be used in conjunction with different pairs of frequencies. The general approach described above works for any two allowable FCC spread spectrum bands. For example, the following combinations are permissible for license-free radio under FCC regulation Part 15.247:

915 megahertz translated to 2440 megahertz.

915 megahertz translated to 5780 megahertz.

2440 megahertz translated to 915 megahertz.

5780 megahertz translated to 915 megahertz.

5780 megahertz translated to 2440 megahertz.

Spread spectrum operation, however, is not required; two licensed narrow bands can be used. However, spread spectrum operation in the 2440 and 5780 megahertz bands is assumed for the remainder of the discussion.

Tag With Time Delay

The tag RF circuitry 300, 400 shown in Figures 3 and 4 use frequency division multiple access, i.e., the tag circuitry 300, 400 receives and emits signals on different frequencies. An alternative embodiment 1500 uses time division multiple access, as shown in Figure 13. For illustrative purposes, assume that the tag circuitry 1500 shown in Figure 13 takes as an input at a receive antenna 1501 a signal at one frequency, such as 915MHz, and emits the same signal through a transmit antenna 1508 at the same frequency after a delay of a microsecond. Assume that a cell controller, such as cell controller 102a, transmits an interrogation signal 106 in bursts every 2 microseconds. A tag, such as tag 101a, takes this signal as an input through the receive antenna 1501. The signal then passes through elements 1502-1504, as in Figures 3 and 4. A time delay element 1505 is then used to delay for a microsecond. The signal then passes through a transmit bandpass filter 1507 and is emitted from the transmit antenna. A SAW device can be used as the time delay element 1505. During the period of the delay, the cell controller ceases transmission, and reflections of the interrogation signal 106 in the environment die down to a minimal level. This half-duplex approach allows single frequency operation, although with lower bandwidth than with a full-duplex frequency shifting approach. As in the frequency-changing tag, the delay-based tag is capable of modulating the response signal by a 180-degree phase shift. In other respects, the tag design 1500 shown in Figure 13 is similar to those shown in Figures 3 and 4.

Cell Controller RF Design

Figure 6 shows the radio stage of a cell controller 102a. The architecture of an antenna module, such as the antenna module 104, is shown in Figure 7. Together, the cell controller 102a and its remote antenna modules 104a-d modulate a baseband square wave input onto a 2440-megahertz carrier, filter the resulting 2440 megahertz signal to comply with FCC transmission requirements, transmit that filtered 2440-megahertz signal through a selected antenna module, receive a returning 5780-megahertz tag response through the same antenna module, extract the I (Inphase) and Q (Quadrature) components of the demodulated baseband signal, and digitize the result for subsequent processing.

Figure 10 shows the main components of the cell controller digital subsystem 650. In summary, the digital subsystem 650 provides a baseband input signal 601, and some number of nanoseconds later receives a demodulated response 107 from a tag 102a. The

microprocessor 1001, as noted above, can change the behavior of the radio system by (a) modifying the baseband input signal 601; (b) modifying the chip rate, pseudonoise sequence length, and/or the pseudonoise sequence code; (c) modifying the transmit frequency 610 of radio transmitter 1002 and the receive frequency of radio receiver 1003 within a narrow range; (d) modifying the transmit gain of radio transmitter 1002 and the receive gain of radio receiver 1003; and (e) by switching antenna modules 104a-d.

The demodulated response 107 from the tag 102a is split into I (Inphase) and Q (Quadrature) components by the receiver Radio 1003, and digitized by a digitizer 636. Signal processing hardware 1004, for example a combination of DSP and FPGA components, reduces the output from the digitizer 636, performing correlation operations at high speed. If binary phase-shift keyed (BPSK) modulation is used on the transmitting side, the I and Q channels are correlated separately and combined. For quadrature phase-shift keyed (QPSK) modulation, each channel must be correlated twice, once with each sequence. The correlated data from the signal processing hardware 1004 is processed by a microprocessor 1001, such as a Pentium processor. Communications between the microprocessor 1001 and the host computer 105 is accomplished using a TCP/IP protocol, with Ethernet being preferred.

The data that is input to the transmit chain is a baseband input signal 601 which is a pseudonoise spreading sequence. The length of the sequence and the code encoded in the sequence are set by a cell controller microprocessor 1001, and can be varied depending on signal processing requirements. Thirty-one or 127 bit sequences are typical, giving about 15 dB and 20 dB of compression gain, respectively. The 2440 megahertz and 5780 megahertz bands can support a 40 megahertz baseband input signal 601, and the cell controller 102a is designed to enable this full bandwidth to be utilized.

Figures 5A-5G show an interrogation signal 106 as it passes through various stages of the cell controller RF circuitry 600. Figure 5A shows a square wave baseband input to the modulator 500. Figure 5B shows this baseband input digitally correlated 510. Figure 5C shows an output 520 from a-modulator 602, viewed through a spectrum analyzer centered at 2440 megahertz. Figure 5D shows a spectrum analyzer view 530 of the tag signal 107, centered at 5780 megahertz. Figure 5E shows the demodulated response from tag 107, separated into its I (Inphase) 545 and Q (Quadrature) 540 components. Figure 5F shows the I and Q components, digitally correlated 550. Figure 5G shows the negative of the second derivative of the correlated waveform, combining the I and Q components 560.

The modulator 602 (Figure 6) modulates the baseband input 601 onto a 2440 megahertz carrier. Various forms of modulation are available and well-known to those skilled in the art. For BPSK modulation, the modulator 602 is implemented as a single pole double throw RF switch 801 that feeds a 180-degree hybrid combiner 803, as shown in Figure 8. The modulator 602 is preferably implemented as a QPSK modulator, which duplicates the BPSK modulator with one channel offset by 90 degrees from the other, each channel driven by a different baseband sequence with acceptable cross-correlation characteristics. Higher order modulation is also possible. Modulation by the modulator 602 results in sidelobes extending for hundreds of megahertz, which need to be filtered to comply with FCC requirements. The 2440 megahertz band has an adjacent band, which imposes very strong filtering requirements, best addressed using in the illustrated embodiment using a SAW filter 607 that combines wide passbands with tight stopbands. A wider passband supports a faster chipping rate in the baseband input signal 601, but a narrower passband provides an opportunity to use a wider range of frequency diversity to avoid jammers and/or support advanced signal processing techniques. The modulator 602 must operate at the same frequency as the available IF filter 607, typically in the range of 200 megahertz to 400 megahertz. A preamplifier 606 is necessary prior to the SAW IF filter 607, and the output of the filter needs to be amplified by an amplifier 608.

A transmit IF oscillator 605, like all other RF oscillators in the cell controller circuitry 600, is phase locked to a 10 megahertz crystal source 603, which is distributed through a filter and splitter network 604 to each of the oscillators. The 10 megahertz source 603 needs to be within a few kilohertz of the 10 megahertz sources on the tags to avoid excessive baseband phase shift.

The output from the IF filter 607 (from amplifier 608) is then mixed by a mixer 609 with the output from a phase locked oscillator (PLO) 611 and is converted to a carrier frequency of 2440 megahertz. The frequency of the PLO 611 can be modified within a narrow range under microprocessor control 610, in order to provide the frequency diversity needed to avoid jammers and/or for various advanced signal processing techniques. The degree of frequency diversity available is related to the specifications of the IF filter 607, with narrower filters permitting a slower chip rate but having more frequency flexibility. Not shown in Figure 6 is a filter typically needed to remove undesirable harmonic and difference frequencies from the output of the mixer 609.

Following the mixer 609 is a driver amplifier 612 which raises the power level of the signal 106, so that it can be driven down the cable 103a to the remote antenna module 104a, and which buffers the output of the mixer 609 for a bandpass filter 613. The RF bandpass filter 613 is needed to remove FCC non-compliant outputs from the mixer 609. A 5 directional coupler 616 provides a port to examine the signal 106 before it is transmitted to the remote antenna modules, for example antenna modules 104a-d.

An attenuator 614 under microprocessor control 615 allows the signal processing software to decrease output power when a tag 101a-d is known to be nearby. This is helpful in circumstances when a nearby tag is known to be over-driven by the cell controller, and/or 10 the signal processing software needs the tag to operate in a more linear range.

The signal is then fed into a diplexer 618, which combines the transmitted 106 and received 107 signals onto a single wire. The diplexer 618 is a highpass/lowpass filter combination 619a that attenuates a received signal 107 relative to the transmit side and attenuates the transmit signal 106 relative to the receive side. Because of the presence of 15 the Tx and Rx bandpass filters 613, 624, the specifications of the diplexer 618 are not very stringent.

The cell controller RF stage 600 shown in Figure 6 supports one remote antenna module 104a-d at a time. In order to support multiple antennas from the same cell controller, the system needs a switch 619, which enables a microprocessor control 620 to rapidly switch from one antenna to the next. The switch 619 takes RF and passes it into one 20 of n cables, where n is, for example either 8 or 16. The switch 619 also provides DC power to the selected line. The RF signal is coupled into the cables with a capacitor (not shown), which provides DC isolation, and the DC power is coupled into the cables with RF chokes to provide RF isolation. Thus, the DC and RF travel together through a single coaxial cable 25 to the selected antenna.

The rise time of the DC in an antenna is in the range of 20 microseconds, limited by the effective resistance of the circuitry in and characteristics of the antenna and the capacitors needed for operation. To provide antenna switching time in the microsecond range, the DC power to an antenna is preloaded before the RF is switched.

30 Referring to Figure 7, in an antenna system 700, the combined DC and RF signals arrive through a coaxial cable, such as the cable 103a from the cell controller 102a. A bias-tee 701 separates the RF signal 710 from the DC signal 712. The DC signal 712 is sent to Tx/Rx power control logic 702 which, in the simplest embodiment is a filter to remove

noise from the line and provide a clean 5 volt power source. The RF output 710 from the bias tee 701 is fed into a diplexer 715, which is identical to the diplexer 618 in the cell controller 102a. This is then amplified by an amplifier 703 to the power level allowed by the FCC, and filtered by a filter 704 to remove line and amplifier noise in compliance with FCC regulations. The resulting signal is then sent to a transmit antenna 705.

The transmit antenna 705 and receive antenna 706 are, in this embodiment, wall mounted patch arrays, providing reduced energy in the vertical direction and spreading energy laterally, so that power is not wasted in the floor and ceiling, and so that minimal power is radiated upward. The 5780-megahertz response 107 from the tag 101a is filtered by a filter 707, amplified by an amplifier 708, and sent back down the cable 103a to the cell controller 102a.

The system is designed to use cables 103a-d of a standard length, for example, 20 meters. A cable extender module 900 connects two lengths of cable and supports an extended cable length. Referring to Figure 9, the elements of the module 900 use the DC power 910 from the cable 103 to drive low noise amplifiers 903, 904, which provide enough gain to drive the next section of cable. Bias tees 906, 907 separate the DC power 910 from the RF signals, and dipoles 908, 909 operate to separate the transmit signal 106 from the receive signal 107.

Referring to Figure 6, the signal returning from the antenna module 104a to the cell controller 102a passes through the switch elements 621, 619 and diplexer 618 to the cell controller receive RF chain 622. The signal passes through a combination of a preamplifier 623 and bandpass filter 624, the exact arrangement of which varies based on the parts selected. A digitally controlled receive attenuator 625 under microprocessor control 626 is used to avoid saturating the receive chain when the tag 101a is known to be nearby. This is necessary to avoid losing the relationship between the I and Q components of the received signal 107, necessary for proper correlation and other signal processing.

The signal then enters an I-Q zero IF demodulator circuitry 627-633. As noted previously, the microprocessor Rx frequency control 635 must be set in tandem with its counterpart in the transmit chain. The resulting signal, an example of which is shown in Figure 5E, is fed into a digitizer 636 (Figure 10) in preparation for digital signal processing.

The embodiment described above is simplified, based on an assumption that the cell controller can send and receive from only one antenna at a given time. Improved performance can be achieved by selecting send and receive antennas independently of each

other. Software in the cell controller determines which antenna module receives the best signal from the tag. For example, if a particular tag, such as tag 101a, is close to an antenna, such as antenna 104a, then the antenna 104 will receive a strong signal from the tag 101a. The cell controller 102a can then transmit a signal, such as signal 106, from antenna 104a, and receive the transponded response 107 at antennas 104b, 104c, and 104d in turn. This can result in a stronger signal received at the antennas 104b-d, compared to the signal the antennas 104b-d would have received had the signal 106 been transmitted and the signal 107 received from each of the antenna modules 104b-d independently.

The design 1600 shown in Figure 14 provides multiple receive chains 1610a-1610n operating in parallel. Each of the receive chains 1610a-1610n includes an IQ demodulator, a digitizer, and a correlating element, shown as integer DSPs, for example, integer DSP 1620. Implementing each receive chain on a separate card provides scalability. The use of multiple receive antenna modules for the same transmitted signal allows the cell controller signal processing software to employ spatial processing techniques to isolate multipath effects. These techniques take advantage of the fact that the multipath-corrupted response will have different characteristics at each antenna.

Bit Detection

In an ideal environment, a simple triangular correlation peak can be derived from a received tag signal 107, as shown in Figure 5B. Distortions introduced in the radio chain, particularly caused by indoor multipath effects, result in a distorted but nonetheless distinct correlation peak, a function of which is shown in Figure 5G. For the purpose of bit detection, the essential point is to reliably detect the existence of a series of correlations, which indicates the operation of a tag. Figure 11 shows how tag data is extracted from a series of correlations. In the left half of the chart 1110 shown in Figure 11, the tag is transmitting a "zero." This is accomplished by setting the tag's modulator 307 to pass the interrogator signal 106 unaltered. When the received tag signal 107 is correlated with the transmitted pseudonoise sequence, essentially identical correlation peaks result. Three such peaks 1120a-c are illustrated here. During the time of the fourth correlation 1120d, the tag flips the phase of the modulator by 180 degrees, indicating a "one," as shown in the chart 1110. Since the modulation is changed in the middle of a bit, the fourth correlation data peak 1120d is corrupted, and is best ignored. The fifth and sixth correlation peaks 1120e-f cleanly reflect the 180-degree shift.

Pseudonoise sequences can be varied under microprocessor control at the cell controller. When a tag's presence is first detected, relatively short sequences must be used, as shown in Figure 11. Once the tag's bit timing is ascertained, it is possible to use longer sequences for improved SNR, which is helpful in distance measurement.

5 An important consideration, not shown in Figure 11, is that the balance between the Inphase and Quadrature components of the received signal will drift over time. This is because the 10-megahertz sources in the cell controller 102a and the tag 101a will differ typically by a few kilohertz. This factor can be calibrated by noting the phase difference between contiguous correlations, detectable at the baseband by noting changes in Inphase
10 and/or Quadrature components of the received signal. As noted previously, this same calibration process can be used to calibrate the tag clock in reference to the cell controller clock, allowing for precise prediction of tag chirping time, without needing to precisely measure the timing of tag bit transitions.

15 Interaction Between Cell Controllers and Tags

Each tag is a stand-alone unit that is unaware in any way of the outside world. Each tag has a Unique Identifying Code (UID) associated with the tag when it is manufactured. A tag wakes up periodically and, for a short period of time, converts any incoming 2440-megahertz signal 106 to an outgoing 5780-megahertz signal 107, while modulating its UID
20 and other data onto the outgoing signal 107 which it chirps (transmits). The tag does not communicate with other tags. The tag does not explicitly respond to an interrogation signal, but merely transponds any incoming signal 106 in the 2440-megahertz band, which may or may not include a pseudonoise sequence from a nearby cell controller antenna module 104a. This approach greatly simplifies the design and fabrication of the tag 101a.

25 Some portion of the time, two or more tags will transpond simultaneously. In many cases, one of the two tags will return a stronger signal than the other tag and some data will be lost in such a collision. To avoid collisions occurring in a repeating pattern, tags "wake up" and chirp their UIDs at randomized times, which can be calculated (by both the tag and the cell controller) based on a pseudorandom number generator which incorporates the tag's
30 UID. For example, for a tag which chirps approximately every 5 seconds, the tag generates pseudorandom numbers between 0.0 and 2.0, and adds these to a 4.0-second minimum delay time, resulting in a sequence of delay times uniformly distributed between 4.0 and 6.0 seconds.

It is possible to use analog inputs, such as an input from an internal clock or the delay of an RC circuit, to the pseudorandom number generator to vary the seed over time, but a purely digital scheme is preferred to enable the cell controller 102a to accurately forecast a known tag's chirping time. A typical pseudorandom number generator has the form:

$$N = \text{rand}(\text{seed}) \quad \text{Formula 1}$$

The resulting N is used as the seed for the next pseudorandom number in the pseudorandom number sequence. When using a pseudorandom number generator of this kind, it is possible that two tags will use the same seed, resulting in their tag signals repetitively colliding. Further, with small differences in tag clocks, all pairs of tags will eventually drift through this synchronized state for some amount of time. To avoid these conditions it is desirable to incorporate each tag's UID, as noted above, into the delay time for that tag, resulting in a different pseudorandom sequence for each tag, that is:

$$\text{Delay} = f(N, \text{UID}) \quad \text{Formula 2}$$

One simple example of such a function is:

$$\text{Delay} = \text{Xor}(N, \text{BitRotate}(\text{UID}, \text{AND}(N, 1111_2))) \quad \text{Formula 3}$$

Referring to Formula 3, it is possible to reconstruct seed from UID, Delay and $\text{And}(N, 1111_2)$, by calculating $N = \text{Xor}(\text{Delay}, \text{BitRotate}(\text{UID}, \text{And}(N, 1111_2)))$.

Referring to Figure 12a, one embodiment of the tag datagram 1400 contains a header 1401 to enable the cell controller to detect the tag's presence, followed by an identifier preamble 1402, followed by the tag's UID 1403. The header 1401 can be of zero length. The identifier preamble 1402 can be implemented, for example, as a validity check such as a cyclic redundancy check (CRC). Given a sufficiently simple Delay function and high clock stability, the cell controller can infer the tag's chirping sequence by noting the timing of a series of chirps of the datagram 1400.

Referring to Figure 12b, in another embodiment of the tag datagram 1410, the tag adds Delay information 1414, thus enabling the cell controller to forecast the transmission time of the tag's next and subsequent chirps of the datagram 1410. In the example of Formula 3, this information would include the data: Delay and $\text{And}(N, 1111_2)$.

Referring to Figure 12c, in another embodiment of the tag datagram 1420, a shorter header is used than in the datagrams 1400, 1410 of Figures 12a and 12b, such that the cell controller is not guaranteed to have enough time to detect the tag's presence before the UID 1423 contained in the tag datagram 1420 is transmitted. Appended to the datagram 1420 is

the transmission delay 1425 of the next chirp, enabling the cell controller to anticipate the time that the tag will next chirp its datagram 1420, even if the cell controller does not have enough time to identify the identity of the tag from the first received chirp of the datagram 1420. The cell controller can then anticipate this next chirp and ascertain the identity of the tag at that time. Once the tag is identified, the cell controller can duplicate the tag's pseudorandom number generator to calculate the times of all future chirps by the tag. In the tag datagram 1420 of Figure 12c, a series of special synchronization bits 1424 are inserted between the UID 1423 and the delay information 1425, to reliably determine when the UID 1423 ends; in this case, the UID 1423 must be defined so that it does not include the synchronization sequence or its inverse.

Figures 12a, 12b, and 12c include optional data sections 1404, 1415, 1426, which allow a tag to transmit data to the cell controller. These sections 1404, 1415, 1426 can include data from within the tag, such as from a motion detector or a low power indicator, or from an external device attached to the tag, such as metabolic information for medical telemetry if the tag is attached to a person.

An identifier preamble, related to the tag UID, precedes the tag UID. This identifier preamble enables the cell controller to quickly verify that a tag is chirping as expected, without needing to decode the tag's complete UID. This frees the cell controller for other activities, such as communicating with different tags in proximity to other antennas. The identifier preamble 1402, 1412, 1422 and tag UIDs 1403, 1413, 1423 are set externally, and if appropriate can be defined to include error correction bits.

The UID of a tag may be hardcoded into the tag (e.g., as a serial number). Tags may be grouped based on their UIDs, and different groups may be associated with different cell controllers. Each cell controller contains information (received from another source) about which tags are in the group associated with the cell controller. When a cell controller receives a tag signal, the cell controller can extract the UID information from the tag signal to determine whether the tag signal was sent by a tag in the group associated with the cell controller.

In the tag datagrams 1400, 1410, 1420 of Figures 12a-c, the delay information fields 1414, 1425 and data fields 1404, 1415, 1426 can also include error correction bits. For simplified processing, data can be reduced to a stream of half-bytes. To determine what value to send for a particular half-byte, the tag can look up the half-byte's value in a table which contains, for example, 8-bit values, which represent the value of the half-byte plus

error correction information. A single cell controller can handle all three types of datagrams 1400, 1410, 1420 shown in Figures 12a-c. The choice of datagram type would depend on the application requirements for a particular tag.

The amount of time it takes for a cell controller to detect the presence of a tag may
5 vary depending on the nature of the cell controller design. For example, a 100-microsecond time to switch antennas may be significant when the cell controller is cycling among 16 antennas. In order to be assured that a tag will be identified the first time its tag signal is received by the cell controller, the tag datagram header must be long enough to give the cell controller time to try all of its antennas. If the performance requirement is in the range of
10 100 tags per second, 2 or 3 extra milliseconds in the header can be tolerated. But for higher performance requirements, or when tag power consumption must be minimized, it is necessary to either improve the performance of the cell controller or to use a tag datagram 1420 of the type shown in Figure 12c.

By anticipating time of transmission from a particular tag, the cell controller can
15 collect tag information from a variety of antennas in an organized way, in order to better calculate tag location by using antenna and/or frequency diversity. If a tag is responding exactly when it is expected to respond, it is not necessary for a cell controller to detect every bit transmitted in the tag datagram in order to be reasonably certain that it is receiving a signal from the correct tag. A correct identifier preamble arriving exactly on schedule is
20 almost certain to be from the expected tag. This provides an opportunity for the cell controller to try a variety of antennas that may or may not be able to communicate with the tag.

If it is necessary to track tags between the tag datagram transmissions, the tag can be configured to send shorter transmissions more frequently. For example, if a tag is
25 configured to chirp its datagram on average every 10 seconds, it may also be configured to transmit a much shorter code more frequently, such as every half-second. This shorter code might be as short as one bit long, and take just a few microseconds to transmit. Thus, even hundreds of such transmissions per second would consume only a small percentage of the communication channel. The cell controller can anticipate the exact timing of each such
30 transmission, thus matching each signal with the originating tag based on the time of the transmission. Error correction codes can be arranged such that long chirps from one tag will not typically be corrupted by quick chirps from other tags. The cell controller has the data to forecast most of such collisions.

When a tag is first introduced into a cell controller's range, collisions between datagram transmissions of different tags will occur in an unpredictable way. Transmission by a tag newly introduced into the area, or by a tag which spontaneously increases its transmission rate (such as in response to a motion detector or "panic button"), cannot be forecast by the cell controller 102a, and might cause data corruption. However, once the tag is identified, previous collisions can be modeled and questionable data discarded.

Alternatively, if signals from two tags collide, the cell controller can select an antenna such that the signal the antenna receives from one tag will be stronger than the signal the antenna receives from the other tag.

In a more advanced tag design, the cell controller has a means for sending information and instructions to the tag during the times that the tag is known by the cell controller to be in operation. Such instructions can include commands to be passed on to a device attached to the tag. The cell controller is capable of downloading such information, most simply by on-off keying, or for more advanced tags by flipping the pseudonoise bit sequences to indicate one or zero. In general, the downlink (downloading) approach is driven by tag cost and feature requirements, with higher bit downlink transmission rates requiring more expensive receivers that consume more power. Thus, a single cell controller can support read-only tags, read/write tags, and high-speed read/write tags simultaneously, with the cell controller adapting its behavior depending on the features supported in a particular tag.

Timing of transmission from tag to cell controller depends on the item being tagged. Inventory and equipment can be set to transmit relatively infrequently, such as once per minute. More frequent transmissions would be required, for example, for tags on people in a secure facility. For read/write versions of tags, timing of transmission could be modified on command from the cell controller.

Alternative tag designs enable variation in transmission time based on environmental factors. For example, motion detectors may be placed in a tag to decrease the time between transmissions when the tag is in motion. As another example, a tag might transmit more frequently and with higher power when the tag has been tampered with. As another example, a tag might incorporate a slightly modified Electronic Article Surveillance (EAS) device, which would cause the tag to transmit its UID more frequently when in range of a standard EAS detector. More generally, if a tag is attached to another electronic device, transmission interval can be modified under the control of that device.

Tag Power

Tags 101a-c transmit a low RF power level in order to increase their portability and lifespan. In addition, tag signal transmissions 107 are designed to be only a few milliseconds in duration. Therefore, even if a tag transponds its UID every few seconds, careful tag design allows the battery life of the tag to approximate the shelf life of the battery itself. For even lower power usage, a motion detector can be incorporated into the tag so that, for example, transmissions can be less frequent when the tag is stationary.

For some situations, battery replacement may be accomplished by incorporating the battery in an attachment mechanism. For example, re-usable tag electronics may be attached to a disposable patient bracelet, with the battery included in the bracelet. As another example, a battery may be incorporated into the clip of an ID bracelet. More generally, a battery may be incorporated into an inexpensive disposable portion of an active RFID tag, with the electronics in the other, more expensive, portion.

If the item to which a tag is attached is a power source itself, the tag could tap into that power source. This approach is most practical in situations where the tag can be designed into the equipment itself (such as a handheld computer), or where the equipment and power source are large (such as a forklift). A larger power source allows for longer tag range.

Estimating Tag Location

A tag signal 107 is received at a time that is the sum of (1) known fixed delays in the cell controller 102a that transmitted the interrogator signal 106, due to its circuitry and the wiring to and from its antenna modules 104a-d, (2) fixed time delays in the antenna module 104a and tag 101a, and (3) the time for the interrogator signal 106 and tag signal 107 to travel through the air.

Since (1) and (2) are fixed, attention can now be turned to (3), the travel time of the interrogator signal 106 and tag signal 107 through the air. The duration of the pseudonoise sequence modulated onto the 2440MHz carrier signal 106 by the cell controller 102a must be greater than the combined travel time of the signal 106 and the tag signal 107.

Techniques for correlating a pseudonoise sequence are well known in the art. In the absence of multipath effects, the cell controller 102a can derive a simple triangular correlation peak from the received tag signal 107, as shown in Figure 5B. But in most indoor environments, the actual received tag signal looks more like that shown in Figures

5D-5G. Indoor radio signals are subject to substantial multipath effects, due to reflection from a variety of surfaces, such as whiteboards, fluorescent lights, file cabinets, elevator shafts, steel beams, and the like. When a tag 101a transmits a tag signal 107, the sum of the direct tag signal 107 and reflected signals is received at the cell controller antenna module 104a. A variety of approaches can be employed in order to extract a correlation peak from such information, with the particular approach chosen depending on the available signal quality, processing power, and required performance.

A 40-megahertz chip rate results in a correlation peak with a rise time of 25 nanoseconds, corresponding to a rise time distance of about 25 feet. Since tag location is calculated using round-trip travel time, single-chip accuracy therefore allows tag distance to be calculated within about 12 feet, without any advanced signal processing.

Approximate location of a tag can be calculated by noting when correlated signal-to-noise ratio rises above a predefined level. Improved accuracy can be achieved by trying a small variety of carrier frequencies and choosing the one that rises the earliest; such frequency diversity is supported by the radio system shown in Figures 4-6. This approach is sensitive to the system signal-to-noise ratio.

An alternative approach finds the peak of the correlation function. For an improved result, the signal delay is measured by taking the negative of the cross-correlation function's second derivative and finding the location of its peak, as shown in Figure 5G.

For highest accuracy, the MUSIC algorithm, known to those in this field, can be used, for which accuracy in the 0.01 chip range has been reported. MUSIC requires frequency diversity, which is supported by the radio system herein disclosed and shown in Figure 6. The method is based on a decomposition of the eigenvector space of the pseudonoise correlation matrix of the delay profile data vector. Frequency diversity is required, where each distinct frequency provides the information to solve for an additional multipath component. For tags that are mostly stationary, necessary data can be collected and the calculation completed as a background process. For inventory applications, motion detectors can be incorporated into the tag, which would then inform the cell controllers whenever their locations need to be recalculated.

In situations where adequate data collection time, antenna diversity, frequency diversity, or processing power are not available, a variety of heuristic techniques can be used to calculate tag location, even in the presence of severe multipath effects. A variety of

techniques well-known in the art can be used to estimate bearing to each antenna, which can be employed as required.

In many circumstances, although precise accuracy in calculating tag location is not required, it may still be important to calculate tag location with reference to a floor or partition. Story-by-story discrimination in a building can be accomplished by mounting antennas in the ceiling that radiate downward (or antennas in the floor that radiate upward), as shown in Figure 2. Similarly, antennas mounted sideways on horizontal partitions can determine location relative to those partitions. Relatively narrow beamwidth antennas, which are less sensitive to multipath effects, can be directed at doorways and the like.

A cell controller antenna 104a can be mounted near a computer screen, with coverage corresponding to viewing angle of the screen. Software can then be arranged to automatically configure the operating system for the person in range, or to blank the screen for security purposes depending on who is or is not in range. Similar concepts can be used in conjunction with copying machines, microfilm readers, restricted equipment, and the like.

A single antenna module can include three separate antennas placed in a triangle. By comparing phase difference using the Inphase and Quadrature components of the returning signal, an indication of the tag's angle can be determined. In high frequency embodiments, such as at 2.45 gigahertz, such antennas could be within inches of each other and be quite effective.

Heuristic techniques can be used to analyze the correlation profile to estimate the time at which the correlation began, that is, the time at which the correlation peak begins to be distinguishable from the "noise" baseline. Frequency diversity can provide a variety of samples, the best of which can be chosen. Improved estimates can be achieved by pattern matching the correlation peak with one of a vocabulary of well-studied typical correlation profiles. For calibration purposes, tags can be placed at known fixed locations, and tags passing near those locations will likely demonstrate similar correlation profiles. Such fixed tags can also be utilized to detect the presence of jammers (objects which emit, intentionally or unintentionally, interfering signals) and to provide a real-time testbed for trying a variety of antijamming techniques.

Antenna diversity is the most important tool for improving the accuracy of tag location calculation. If low accuracy is required, antennas can be placed so that only one or two antennas are in range of a given tag. In this case there is insufficient data for triangulation, and only enough information to detect tag presence and estimate tag distance

from one or two antennas at any one time. Approximate bearing of a tag can be estimated from the signal strength of antennas designed for the purpose and well-known to those skilled in the art; with the caveat that such bearing tends to reflect the strongest signal received, which might include a substantial multipath component. Conversely, for areas requiring high accuracy, a diversity of relatively narrow beamwidth antennas can be installed, for example, at entrances, which together provide a clear picture of location.

Antenna diversity also provides system scalability. For other facilities, or sections within facilities, that do not require calculation of tag location with high accuracy, substantially omnidirectional and/or ceiling mounted antennas can be installed relatively far from each other, for a relatively low cost per square foot of coverage. For facilities, or sections within facilities, that require high positional accuracy, a diversity of closely-spaced and/or directional antennas can provide high accuracy at an increased cost.

Cell Controller Operation

The design of the overall system, and the fact that the intervals between tag signal transmissions are generated pseudo-randomly, provide an opportunity for secure operation. Tags dispersed at the entrance of a secure facility can be tracked through that facility and a special code can be emitted when the tag is tampered with. Although the tag's code can be determined by monitoring the tag's response, the tag's transmission interval varies according to an algorithm that can be arranged to be known only to the tag and the host, and impossible to directly determine without destroying the tag. The tag may for example, include an element, such as a physical element, for reprogramming its code and transmission interval. For example, a photo ID with an incorporated tag might be reprogrammed each time the person wearing the photo ID passes a security checkpoint, potentially in conjunction with biometrics technology.

For full coverage of a facility, multiple cell controllers can be installed, covering somewhat overlapping areas. Although each cell controller will be operating according to a search and data collection method, rapid movement between antennas, the pseudonoise code, the changing chipping rate, and so forth, will appear as random noise to another cell controller. In addition, codes with known cross-correlation characteristics, such as Gold Codes, can be allocated to various cell controllers by the host computer, particularly the codes used to search for tags. Alternatively, cell controllers can switch choice of pseudonoise codes on a randomized basis.

For tags on the border between two cell controllers, each cell controller reports the tag's distances from its antenna modules. The central host 105 assembles this data to calculate the tag's location.

A variety of pseudonoise codes are available for use by a cell controller. Thus, if one code seems to be receiving interference from other users of the spectrum, the cell controller can choose another code. The tag, being in essence a transponder, does not need to know the particular code being used. Likewise, center frequency can be adjusted somewhat if another user is causing difficulty.

10 Digital Signal Processing

Referring back to Figure 10, the signal processing hardware 1004 performs operations upon the output of the digitizer 636. The signal processing hardware 1004 generally functions to perform operations upon the signal from digitizer 636 such as, for example, the previously described correlation. Additionally, signal processing hardware 1004 may perform a filtering process to filter out the noise components of received data signals. Various filtering techniques are known to those skilled in the art, such as the use of an adaptive transversal filter. Following is a description of functions that may be performed by the signal processing hardware 1004 in a preferred embodiment of the article tracking system of depicted in Figure 1.

20 Referring now to Figure 15A, shown is a block diagram of a preferred embodiment of the signal processing hardware 1004. Generally, the signal processing hardware may include one or more hardware components that collectively perform digital signal processing of the received signals. Recall that the signal processing hardware 1004 was previously described in connection with Figure 10. The processing set forth below may be implemented using any combination of conventional off-the-shelf hardware and/or software, utilizing DSP and/or FPGA hardware and/or semi-custom or custom ASICs, that may be configured by one of ordinary skill in the art using the description set forth herein.

30 Shown in Figure 15A is signal processing hardware 1004 which is connected to the microprocessor 1001 and the digitizer 636, as previously described in connection with Figure 10. Included in the signal processing hardware is a correlator unit 1800 connected to other DSP hardware components 1802. The number and type of components that are included in the other DSP hardware components 1802 vary with the type of processing done in each particular implementation. Each of the hardware components included in the signal

processing hardware may be controlled by the microprocessor 1001, such as by using connections 1804a and 1804b. The connections required vary with each particular implementation and associated processing.

One embodiment of signal processing hardware 1004 uses only the correlator 1800. In this instance, data is directly output by the correlation unit, as indicated by the path 1804a. Generally, the other components 1802 and associated connections, such as 1804b, would not be used in this embodiment. As will be described in paragraphs that follow in connection with Figure 15B, only a portion of the method steps of Figure 15B are performed by the signal processing hardware in this embodiment. In an alternate embodiment of the signal processing hardware which will also be described in paragraphs that follow, all of the method steps of Figure 15B are performed as part of digital signal processing.

Referring to Figure 15B, shown are method steps of an embodiment for filtering received signals to enable determination of when a tag signal has been received. At step 2000, a number of samples of received signals are taken. The number of samples to be taken varies with each particular implementation. Generally, in accordance with well accepted principles known to those skilled in the art, the sampling rate is typically twice the transmission rate. For example, in one embodiment 254 points or received signals are sampled where each of the 254 signals is provided every 12.5 nanoseconds corresponding to the 80 megahertz sampling rate in which the system has a chipping rate of 40 megahertz. In this system, the signal transmitter rate is 127 chips per bit, or rather, every 3.2 microseconds a bit of data is transmitted. Generally, "chipping rate" and "chip time" are described in "The Practical Engineer", IEEE Spectrum, Vol. 35, No. 9, September 1998.

At step 2002, a first range of the sample received signals is determined. The first range is a subset of the samples recorded, as in step 2000. Generally, certain factors which are dependent upon each particular implementation may be considered when determining the starting point and size of the first range of step 2002. One factor that may be considered is the anticipated arrival time of a signal returned by a tag. This relates to, for example, anticipated delays in the transmission circuitry and may be used in determining a starting point of the first range.

Another factor that may be considered relates to the tag transmission range and the distance at which tags may be expected to be located. This may affect both the starting point of the first range as well as the size of the first range. For example, if objects in one

system are known to be located within a small range, then the earliest possible time which a signal may be received by an antenna of a cell controller is earlier than a time of a different tag in another system in which objects are known to be located farther away from the transmission source. This may affect how many of the received signals which occur earlier in the sampling may be disregarded.

Generally, the data collection event, the transmission of the signals, and the anticipated hardware delays and other timing delays may be calibrated in accordance with the starting point of the collection or sampling. In one particular embodiment, due to the calibration of the transmitting and the data collection as well as the anticipated delays, approximately the 70th sample was determined to be the beginning of the first range because this was one of the earliest points at which a return signal may be expected in a particular embodiment. Also in this embodiment, the actual span or size of the first range is 26 which was determined in accordance with the anticipated range of distances of objects to be located.

At step 2004 auto-correlation is performed for the first range of samples producing a magnitude for each sample in the first range. The performing of auto-correlation upon received signals was previously described. For example, Figure 5B shows the results of performing auto-correlation upon the wave form 500 as shown in Figure 5A. However, the waveform 510 of Figure 5B does not include any "noise" in addition to the originally transmitted signal 500 as shown in Figure 5A. One preferred embodiment implements the auto-correlation portion using a field programmable gate array (FPGA). The ability of the FPGA to perform massively parallel operations, such as matrix operations, efficiently in hardware is one factor in considering using the FPGA to perform the auto-correlation function of step 2004. Generally, use of particular hardware, such as the FPGA, may reduce the real time calculations and computational costs associated with performing expensive matrix operations, as in the method steps of Figure 15B.

Generally, in the calculation of the autocorrelation function, 254 consecutive samples of received data ("sample waveform") are compared with an idealized version of the same data ("reference waveform"). If a coherent demodulation is available, then a real correlation may be performed. Otherwise, a complex correlation may be performed.

Autocorrelation may generally be defined as the integral:

$$\psi(t) = \int_{-\infty}^{\infty} f(t) f(t - r) dt$$

The above equation is a measure of the similarity between a signal and a phase shifted replica of itself. An autocorrelation function is a plot of autocorrelation over all phase shifts ($t-r$) of the signal, where Δt , the change in time, is in half-chip intervals.

For a 127-chip sequence, sampled once every half chip, each correlation calculation 5 takes 254 multiply-and-add operations, and calculating the entire autocorrelation function takes 254^2 which is approximately 64,000 multiply and add operations. If a complex reference waveform is used, computational complexity is increased by a factor of 4. Even with very fast hardware or specialized signal processing hardware, this number of calculations may cause a "bottleneck"

10 due to the amount of time required to perform the calculations.

A correlator implemented in hardware can generally make a quick estimate of tag location by combining various techniques, some as described above and other which will be described in the paragraphs that follow.

Since the reference waveform generally includes 1's and -1's, 2's complement 15 arithmetic may be used for the multiplication operation. With this simplification, one of ordinary skill in the art may implement the foregoing using a field programmable array (FPGA) and/or semi-custom or custom ASICs, enabling operations to take place in parallel and with generally high throughput.

It is generally not necessary to calculate each interval of the autocorrelation 20 function. Some of the range may be ignored because the tag is low-powered and can only be detected at a limited distance. For example, if the tag's radio has a maximum range of 100 meters, there is no reason to perform the autocorrelation function with phase shifts corresponding to the distances in excess of 100 meters. Additionally, in searching for the leading edge, phase shifts of a full chip or more may be used to search for the signal, and 25 then half-chip intervals may be used in the neighborhood corresponding to the time when the first signal is detected. More generally, a subset of the 254 autocorrelation offsets may be used in the search for the peak or the rising edge of the autocorrelation function. This is described in more detail in paragraphs below.

At step 2006, heuristics are used to select a sample point which approximates where 30 in the first range of received signals is the received tag signal. Generally, step 2006 produces a rough estimate as to the timing of the returned signal. This selected sample point is used, in connection with step 2008, to further refine and limit the sampled data points considered in determining the returned signal. One heuristic or technique which

may be used to approximate the location in the first range of the received tag signal is related to the strength or magnitude of the signals within the first range. By looking for the autocorrelation peak in the first subset, this corresponding signal may be used in approximating where the received tag signal may be located. Generally, this is based upon
5 the premise that the strongest received signal corresponds to the direct path of the received tag signal.

The rising or leading edge detection technique is a second heuristic that may be used to approximate where in the first range of receive signals is the actual receive tag signal.
Generally, the samples are observed until a large or significant change in slope is detected.
10 The actual determination of what is "large" or "significant" is relative to each system and varies with each implementation. One technique used with the rising or leading edge detection may include using a normalized value from 0.0 to 1.0 where 1.0 corresponds to the signal with the maximum amplitude received. When two points are encountered in which the slope of the line formed between these two points is greater than, for example, 20
15 percent of the normalized value, then this change may be considered large enough to signal a significant change in slope. Generally, the rising edge detection technique is based upon the assumption that the first peak is the line-of-sight returned signal. Note that this is different than the premise or assumption of the first technique which is based upon the assumption that the strongest returned signal is the returned tag signal.

20 Yet another heuristic is a threshold detection technique. A threshold value is determined, and the first sample point having a magnitude greater than or equal to this threshold is the selected sample point. The threshold value chosen varies with environmental and implementation. Running trials of the system is one suggested method for choosing a threshold value.

25 At step 2008, a second range of samples is determined which is a subset of the first range of samples using the approximated location as determined in step 2006. In determining the second range of step 2008, the precise starting and end point as well as the size or span of the range must be determined. Both the span of the second range as well as the precise starting and ending points of the second range may be related to or dependent
30 upon the heuristic used to approximate the location of the received tag signal in step 2006. For example, if the peak or maximum amplitude of the received signal were used in determining the approximate location of the receive tag signal, one common technique would be to take a specified number of equal points to the right and to the left of this peak

and use this to correspond to the span and beginning and end points of the second range. If a different technique were used, such as the leading or rising edge detection technique, a varying number of points before and after the rising edge may be used. One preferred embodiment, for example, may use the rising edge detection technique or the threshold detection technique. In this embodiment, a range is determined having a starting point which is three to the left of the rising edge, and eight points to the right of the rising edge.

The second factor to be considered is the actual size of the second range. It should be noted that the size of the second range in one particular embodiment is 12. The reasons and factors that may be considered when choosing the size of the second range will become more apparent in light of following paragraphs describing the different operations which are performed upon the second range of data.

In step 2010, the recursive least squares (RLS) technique using the second range of samples is performed. Generally, the RLS technique is used in the design of adaptive transversal filters and is based upon the least-means square adaptation method, as generally known to those skilled in the art. Functionally, the RLS technique used in step 2010 is used to filter out the noise component of received signals. In this embodiment, the RLS technique is used to filter out the noise components of the second range of sampled receive signals. The precise steps and how the RLS algorithm works are disclosed in the paragraph that follows.

At step 2012, a vector of filtered samples corresponding to the second range of samples being filtered are produced. As previously described, this vector contains values which correspond to filtered received signals.

At step 2014, the approximate peak corresponding to the filtered received tag signal is determined using the values included in the vector. Detail is described in the paragraphs that follow regarding how the peak is approximated using the values included in the vector produced by step 2012.

In step 2016, the tag distance is determined using the time of the filtered received tag signal. The precise details of how to determine tag distance using the time of the filtered received tag signal were previously described based upon the difference between transmission time and receipt time of the tag signal.

One of the functions of the signal processing hardware 1004 is to filter the noise component out of a received signal. Before describing detailed steps of how the RLS technique is used in performing this filtering process, a general description of how the RLS

technique is used in a feedback control system to perform this filtering function is described.

Referring now to Figure 16, shown is an example of an embodiment of a block diagram of a feedback control system which filters out noise components of received signals. Generally, the function of the block diagrams of Figure 16 combined filter out the noise components of a received signal. In this particular embodiment, Figure 16 depicts a feedback control system which includes a transversal filter 2022 and an adaptive weight control mechanism 2024. Generally, the transversal filter 2022 operates upon an input signal u and represented as an input vector $u(n)$ with a varying number of components. The transversal filter produces an output signal $E(n)$ which provides an estimate of the desired response $d(n)$. In this instance, the desired response or signal $d(n)$ is the actual received signal. The estimated signal $E(n)$ produced by the filter is compared with the desired response signal $d(n)$ to produce an error estimation $\alpha(n)$. This estimation error $\alpha(n)$ is the difference between the desired response signal $d(n)$ and the estimated signal $E(n)$ as produced as an output by the transversal filter. This error value $\alpha(n)$ is used as an input and feeds back into the adaptive weight control mechanism 2024. The adaptive weight control mechanism 2024 is a mechanism for performing the adaptive control process by varying certain parameters which feed back into the transversal filter 2022.

Generally, the transversal filter 2022 may also be referred to as a tap delay line filter. Further description and details of the transversal filter is described in paragraphs that follow in connection with Figure 17. In this particular embodiment, the RLS algorithm extends the use of the method of least squares to provide a recursive algorithm for the design of adaptive transversal filter such that, given the least squares estimate of the tap weight vector, $w(n-1)$ of the filter at time $n-1$, the updated estimate of this vector at time n may be computed. The estimated signal $E(n)$ is denoted as $w^H(n-1) u(n)$. The precise notation of this will be described also in paragraphs that follow.

Referring now to Figure 17, shown is an example of one preferred embodiment of a transversal filter 2022. Generally, a transversal filter includes three basic elements: unit delay elements (2031a-2031m), multipliers (2032a-2032m-1), and adders (2033a-2033m-1). In particular, when unit delay operator 2031a operates on input $u(n)$, the resulting output is $u(n)_1$. The role of the multiplier 2032a in the filter is to multiply the tap input 2030a by a filter coefficient referred to as a tap weight denoted $w_0(n)$. It should be noted that the

asterisk in Figure 17 denotes complex conjugation which assumes that the tap inputs, and therefore the tap weights, are all complex values. The combined function of the adders in the filter is to sum the individual multiplier outputs and produce an overall filter component denoted $y(i)$. In this particular embodiment shown in Figure 17, the number of delay elements is shown as $m-1$. This is commonly referred to as the order of the filter. Each of the components 2031b-2031m-1 operates in a manner similar to 2031a. Similarly, multipliers 2032b-2032m-1 operate similarly to the multiplier 2032a. Generally, the input signal $u(n)$, and the tap weights denoted $w(n)$ are represented as vectors with each element of the vectors corresponding to various components. When the various components are summed, they produce an estimated signal denoted $y(i)$ in Figure 17. $u(n)$ denotes the tap input vector at a particular time n . The tap weight vector $w(n)$ defines the tap weight vector at a particular time n .

Generally, the recursive least squares or RLS technique attempts to choose a tap vector $w(n)$ which minimizes the expected squared error. The error is determined as the sum of the differences between the expected or actual signal and the estimated signal output from the filter. The adaptive weight control mechanism 2024 of Figure 17 is used to determine the weighting factor associated with each of the errors at a particular point in time where the error is the difference between the desired response or the actual output signal, and the estimated signal produced by the filter. The use of the weighting factor is intended to ensure that data in the distant past is "forgotten" in order to afford the possibility of following the statistical variations of the observable data when the filter operates in a non-stationary environment. Thus, use of the weighting factor allows additional weight to be given to the error values which are most recent in time, and give less weight to those error values which are earliest in time. The precise use of a weighting factor and how it relates to the measure of the "memory" of the RLS technique will become more apparent in following text.

Generally, the adaptive weight control mechanism 2024 provides a correction factor which is applied to the tap weights upon subsequent processing of data. In other words, a correction factor determined at time n is applied to the tap weight at a time of $n+1$. As known to those skilled in the art, a scaled version of the inner product of the estimation error and the tap input denoted $u(n-k)$ is computed for $k=0,1,2,\dots$ to $m-1$. The result obtained defines this correction factor which is applied to the tap weights. Thus, the error factor

provides for adjustment or correction of the various tap weights. This is the nature of the feed back mechanism of the system 2018 of Figure 16.

The previous descriptions regarding Figures 16 and 17 present a general description of the transversal filter and its corresponding adaptive control mechanism as generally known to those skilled in the art of adaptive filter theories as may be used in digital signal processing, for example.

Now what will be described is the RLS algorithm as generally set forth in the textbook entitled "Adaptive Filter Theory", by Simon Haykin, ©1986 by Prentice Hall, Inc.

Referring now to Figure 18, shown are method steps of an embodiment for performing the recursive least squares (RLS) method. At step 2034, variables are initialized to initial conditions. Generally, this includes variables used in subsequent method steps such as the loop control variable n . Some of the variables used in Figure 18 processing steps are initialized to a set of initial conditions. Specifically, n is initialized to one in the method described in Figure 18. The tap weight vector w and the tap input vector u may be initialized to 0. The inverse of the correlation matrix denoted $P(0)$ may be initialized to the inverse of a small positive constant time the identity matrix denoted I . The recommended choice of the small positive constant λ is that it should be small when compared to 0.01 times the variance of the data sample $u(n)$.

Generally, λ is a small positive constant which may be referred to as the "forgetting factor". Since this constant is between zero and unity, multiplying a variable by λ reduces the magnitude of the variable. Using the RLS technique, a weighted sum of square errors is minimized by choosing a vector of coefficients. In this summation, the squared errors in the "past" are weighted by higher powers of λ that are more "recent" squared errors. Thus, the coefficients produced by the RLS technique is generally chosen with less regard for errors in the past.

There are many "rules of thumb" for choosing λ which may depend upon the variance of the data, the variance of the additive noise, the number of components to be determined, and the rate of change of the underlying system to be estimated. It may generally be helpful to note that the effective number of errors which are not yet forgotten is approximately $1/(1-\lambda)$. A forgetting factor of 1/100, for example, yields roughly 100 significant errors. If the underlying system to be estimated is time invariant, then λ should be set as close as possible to unity. In the instance of a time-varying system, λ should generally be chosen such that:

- 36 -

$$\frac{I}{1-\lambda} > 3 * (\text{the number of coefficients estimated})$$

and such that generally:

$$\frac{I}{1-\lambda} \leq \text{the number of samples for which system is approximately constant.}$$

- 5 The former condition generally reflects the fact that the RLS technique provides a converged solution in 3 to 5 times the number of coefficients for cases in which the additive noise is not too severe. The latter condition reflects the fact that the RLS technique attempts to find, for each sample input, a single vector of coefficients which may be used to approximate the last $1/(1-\lambda)$ outputs of the time-varying system.
- 10 In accordance with principles described in the Haykin textbook, certain factors should be considered when initializing values for use with the RLS method. In particular, regarding a starting value for $P(0)$, a starting value should be chosen which assures the non-singularity of the correlation matrix.

- 15 After initialization, control proceeds to step 236 where a determination is made whether or not the loop control variable n is less than or equal to the number of desired iterations. If the determination is made that n is greater than the number of desired iterations, meaning that execution of the method steps of Figure 18 is complete, control proceeds to step 238 where the method depicted in Figure 18 stops. If a determination is made at step 2036 that n is less than or equal to the number of iterations, then control
20 proceeds to step 2040. At step 2040, the gain denoted $k(n)$ is computed as:

$$k(n) = \frac{\lambda^{-1} P(n-1) u(n)}{1 + \lambda^{-1} u^H(n) P(n-1) u(n)}$$

Control proceeds to step 2042 where the a priori estimation error denoted $\alpha(n)$ is computed. $\alpha(n)$ is computed as:

$$\alpha(n) = d(n) - w^H(n-1) u(n)$$

- 25 Control proceeds to step 2044 where the tap weight vector for a particular instance in time denoted $w(n)$ is computed as:

$$w(n) = w(n-1) + k(n) \alpha^*(n)$$

Control proceeds to step 2046 where the inverse of the correlation matrix denoted $P(n)$ is computed as:

$$30 P(n) = \lambda^{-1} P(n-1) - \lambda^{-1}$$

Control then proceeds to step 2048 where the loop control variable n is incremented by 1. Control proceeds to the top of the loop at step 2036 where a determination again is made whether or not the loop formed by steps 236-248 has been performed the desired number of iterations.

5 It should be noted that the method step 2044 describes the adaptive operation of this method whereby the tap weight vector w is updated by incrementing its old value by an amount equal to the complex conjugate (denoted by an *) of the a priori estimation error $\alpha(n)$ times the time varying gain vector k(n), hence the name "gain vector". The a priori estimation, denoted $\alpha(n)$ represents the a priori estimation error. Generally the a priori estimation error refers to an estimate of the error based on a tap weight vector that was
10 made at time n-1. The constant λ is a value close to 1. $1-\lambda$ "represents a measure of the memory of the algorithm". As previously described, this is a weighted value introduced in a definition of the cost function based on the error at time n. Generally, λ is a positive constant close to but less than 1. When λ equals 1, we have the ordinary method of least
15 squares. The inverse of 1 minus λ , generally speaking, is a measure of the memory of the algorithm. The special case where λ equals 1 corresponds to having infinite memory. As previously described in accordance with the use of a weighting factor being associated with the error at each particular point in time, λ is used in determining the actual weight given to particular values of the error determined at different points in time.

20 It should generally be noted that this algorithm is deemed to be "recursive" for the fact that updating the tap weight vector at time n, a prior value for the tap weight vector at time n-1 is used. This becomes apparent when method step 2044 is examined where the tap weight vector denoted w(n) is computed as being dependent upon the value w(n-1). It should also be noted that in the flowchart of Figure 18, the asterisk denotes the complex
25 conjugate of a number. Additionally, the superscript of H, as depicted in step 2042 when associated with the tap weight vector w, implies that the tap weight vector w has the Hermitian property. Generally, a complex valued matrix such as the tap weight vector w(n) is Hermitian if it is equal to its conjugate transpose, as known to those skilled in the art.

Referring back to Figure 15, step 2010 performs the RLS technique just described
30 using the second range of sample receive signals. Also recall that a note was made that the size or span of the second range may generally be related to use of the RLS technique. The RLS technique performs matrix operations which are generally expensive in terms of computing time and resources. Thus, this expense is often a factor to be considered when

determining the size of the second range in that the size of the second range affects the dimensions of the matrix and hence the number of matrix operations which are performed in computations of the RLS technique for the second range of samples. Additional factors should also be considered when choosing the size of the second range. Generally, if the 5 number of points considered in the second range is too large, and the noise component of the receive signal includes tonal frequencies or jammers, then the model may follow the peaks of the jammer signals rather than properly fit a curve identifying the filtered received signal. Additionally, if the size of the second range is too small, then enough points may not be considered to properly fit the curve. It should generally be noted that the number of 10 points or data samples to be included in the second range varies with the system and implementation.

The RLS technique used in this embodiment assumes that the observed sequence includes a linear combination of a known number of the data sequence. Generally, the RLS technique attempts to find the combining coefficients for this linear combination which best 15 fits the observation. Generally, the RLS technique assumes a wide-sense stationary process. Generally, the RLS technique also presumes the presence of random additive noise which is uncorrelated with the data sequence. With the presence of "white noise", the sum process remains generally wide-sense stationary. However, with the addition of a tonal frequency or jammer frequency, the sum process is no longer wide-sense stationary. To use 20 the RLS technique when the sum process is not wide-sense stationary, such as may be in the application of this embodiment as used indoors, a corrective factor should generally be considered. One corrective technique that may be considered is in choosing the number of coefficients or taps.

In this instance, with the sum process not being wide-sense stationary, care should 25 generally be taken to insure that the number of coefficients are not overspecified when using the RLS technique. In this embodiment, the noise input to the correlator may be modeled as the superimposition of two components: a wide-sense stationary white noise process, and a non-stationary intermittent plurality of tones of unknown frequencies and amplitudes. Based on this model for this particular embodiment, it is found that the number 30 of coefficients or taps should generally span the main peak of the impulse response magnitude. However, the number of taps should generally not span more than this in this particular embodiment.

Generally, the number of taps may vary with embodiment and application. For each application and embodiment, the observed process should be modeled to take into account all factors and a number of taps chosen in accordance with these considerations as described herein.

As just described, factors that may be considered when choosing a size for the second range include consideration of the time complexity regarding computational expense as well as application for real time considerations when performing complex calculations. Additionally, the size of the second range varies with the environment in which this application will be used. If there will be tonal frequencies, such as a microwave oven within an indoor environment, this should be considered when choosing an appropriate value for the second range. In one preferred embodiment, a value of 12 was used for the size of the second range for an indoor article tracking system which included the previously described transmission rate of 127 bits per second in a sample of 254 points of received. This indoor system may possibly have tonal frequencies and other jammer signals since it is an indoor application as well as multipath noise. Thus, a range size of 12 for the second range was used considering these factors.

Generally, the steps describing the RLS technique of Figure 18 may be performed using a different data sample set for each execution or iteration of the loop. In one particular embodiment, one data set of 254 received signals was recorded. Rather than perform the method steps of Figure 18 with a different set of 254 data samples each time the method steps of Figure 18 were performed for a particular iteration, the same data set is used for each iteration. Ideally, the number of iterations of the RLS algorithm should be as large as possible in order to meet the mathematical convergence for simulating when n goes to infinity. However, in one preferred embodiment, the method steps of Figure 18, as described in the loop formed by steps 236 through 248, are performed for 84 iterations. Through experimentation for this particular embodiment, it was determined that this was an optimal value to be used in the tradeoff of real time application, computational complexity, and accuracy of locating an object. For other systems and other applications, this number may vary.

As an output of the RLS method performed in step 2010, a vector of filtered samples is produced in step 2012. Each element in this vector corresponds to a component of a received signal which represents a filtered signal. In other words, each element of the vector corresponds to a filtered signal with the noise portion removed. The vector produced in step

2012 is a vector of filtered signals in which each element of the vector corresponds to an element of the second range. The RLS technique performed in step 2010 removes the noise component and returns a filtered signal.

At step 2014, using the data points included in the vector of step 2012, a technique is applied which approximates the peak corresponding to the filtered received tag signal. A more detailed description of step 2014 is set forth in paragraphs that follow in connection with Figure 20.

Referring now to Figures 19A-19E, shown are sample waveforms of received signals for a number of sample points. In Figure 19A, shown is an example of a correlated received signal with a small amount of "white noise". It should be noted that the waveform of Figure 19A generally does not include a multipath component, as would be seen in the system of the embodiment described herein. Figure 19B shows a waveform with a low degree of indoor multipath components as may be included in a received signal. Figure 19C illustrates a waveform that includes a medium degree of severity of indoor multipath noise in addition to a transmitted signal, and Figure 19D shows a severe amount of indoor multipath noise added to a transmitted signal. Each of the waveforms shown in Figures 19B-D are waveforms that may be received in different environments within which a preferred embodiment of the invention operates. Shown in Figure 19E is a correlated waveform, labeled "correlation", which is also typical as the output waveform after performing step 2004 of Figure 15. The second waveform, labeled "model" in Figure 19E, is an example of a waveform resulting from graphing the 12 points included in the second range of samples which is output from the RLS technique in one preferred embodiment. As previously described, this vector may be produced in step 2012 after performing the RLS technique in step 2010. The peak of the actual waveform, for example, may be approximated in step 2014, as the 4th data point of the "model" waveform of Figure 19E.

In this particular embodiment, the transmission rate is 127 chips per data bit. In other words, the bit pattern of the transmitted signal repeats itself every 127 bits. A lesser number bit sequence such as 31 chips, may also be used in a particular implementation which requires less processing time and less number of samples to be taken. However, the gain is lost as you decrease the number of patterns in the bit sequence. A longer bit sequence generally allows for a greater distribution of energy over a larger period of time. A longer sequence generally enables a stronger signal at the output of the correlator than

that produced by a shorter sequence. The processing gain indicates this signal strength enhancement and refers to the number of bits in the sequence.

Other factors regarding a bit sequence which should be considered when implementing the techniques described herein relate to the properties of the bit sequence regarding auto-correlation and cross-correlation. For example, a maximal sequence has good auto-correlation but bad cross-correlation. Thus, a transmitted signal which is a maximal sequence would produce an idealized peak waveform, for example as shown in Figure 5A. However, the cross-correlation of such a sequence may not be desirable in an embodiment where better cross-correlation is required. For example, in an embodiment where one is required to detect or reduce the interference between multiple transmitted signals, good cross-correlation is often needed. For applications such as this where one is required to distinguish between multiply transmitted signals, a sequence such as the Gold code may also be used. It should be noted that the precise sequence length as well as the various properties of the sequence sent in the transmission signals may vary with application and each particular implementation.

In the previously described embodiment, a local positioning system was described in which the assumption is made that the articles or persons being tracked by the system may be at a different position at any particular point in time. In other words, there is no assumption that the person or object will remain stationary for a majority of time. In an application of a system which includes the local positioning system and the techniques described herein, if the objects being located are primarily stationary for a majority of the time, then special processing may be performed in the atypical case when an object is determined or sensed to have moved, as by a motion detector. When motion of a primarily stationary device has been detected, special processing may be performed using the techniques previously described. Since it is the atypical case in which the location of an object is to be determined, in the small number of instances when the location of an object or person needs to be determined, the system may devote additional processing time to locating the object. Thus, in an application of article tracking in which the objects tend to remain stationary, minimizing the amount of computational time to the extent as previously described may not be a factor in selecting, for example, the size or span of the second range.

It should be noted that various portions of the signal processing hardware 1004 may be implemented in varying combinations of software and/or hardware dependent upon the particular application and system implemented. For example, in one preferred embodiment,

the auto-correlation function performed in step 2004 on a first range of samples is implemented in a FPGA, as previously described. Dependent upon each particular system, application, and requirements of each system, other functions of the signal processing hardware 1004 as described herein may be implemented in varying combinations of software and/or hardware.

In one of the previously described embodiments, the same data set was used when performing multiple iterations of the RLS algorithm of Figure 18. As an alternative to using the same data set for each iteration, a single data set set may be “reused” on subsequent iterations with slight modifications, such as rotating or shifting to the right or to the left by one data element for each iteration.

Referring now to Figure 20, shown are example steps of a method for approximating the peak corresponding to the filtered receive tag signal using the values included in the output signal vector produced as a result of step 2014. The method steps shown in Figure 20 are more detailed steps of the method step previously described for step 2014 of Figure 15. For all of the magnitudes of the signals included in the output signal vector, determine the largest magnitude, as in step 2100. It should be noted that in one embodiment, if the first or last element of the vector has the largest magnitude, then an assumption is made that the received signal cannot reliably be determined, and the method stops execution.

Subsequently, in step 2102, a threshold value is determined which, in one embodiment is equal to 62.5 percent of this largest magnitude previously determined in step 2100. At step 2104, it is then determined which element in the output signal vector is first in time to exceed the threshold value . This signal determined to be first in time to exceed the established threshold is referred to as the X. It should be noted that in one preferred embodiment, if the “X+1th” element of the vector is smaller than the magnitude of the “Xth” element, then the “X-1st” element is used rather than the “Xth” element.

At step 2106, other elements of the vector are selected to be used in subsequent processing steps to determine the actual received signal. In one embodiment, two other elements are generally chosen. These are the vector elements denoted by the indices “X+1” and “X+2”. It should be noted that the element indicated by the index value “X+2” is not used if the the “Xth” element is the second to last vector element. In step 2108, using a weighted average formula of these 3 points, the actual received tag signal is determined based on the expected shape of the received signal as being a correlated signal with a triangular peak. One embodiment calculates or estimates the received tag signal as

indicated in the pseudo-code type description below. It should be noted that the description below generally summarizes that which is set forth and previously described in conjunction with the method steps of Figure 20.

```
5      Total = 0
      FOR each element in the vector, v, DO
      {
        Determine the weight of current vector element, denoted by index=j, as:
        weight v[j] = MAX(0, magnitude(vector element j) -
                           (largest magnitude of all
                           vector elements/4) )
10     Total = total + weight v[j]
      }
      Received signal = X+1 - (weightv[X]/Total) + (weight v[X+2]/Total)
      or
15     Received signal = [X+1 - (weightv[X]/Total) /* if element X+2 is not
                           used and is the second to last
                           element of the vector **/
```

The following points are worth noting regarding the previously described embodiments of the signal processing hardware 1004 and previously described method steps of Figure 15B. In one of the previously described embodiments of the signal processing hardware which includes only the correlator 1800, method steps 2008-2014 of Figure 15B are not performed. Rather, only method steps 2000-2006, and 2016 of Figure 15B are performed. Specifically, in this embodiment the autocorrelation (step 2004) is performed by the correlator 1800. The approximation produced as a result of using the heuristics at step 2006 is considered to be the received tag signal. A first heuristic, as used in step 2006, determines the received tag signal to be the signal in the first range with the maximum magnitude of all the signals in the first range. A second or alternate heuristic that may be included in an embodiment (at step 2006) is to choose a threshold value. The first signal included in the first range having a magnitude equal to or exceeding this magnitude is determined to be the received tag signal. The method for choosing the threshold value may vary with environmental and other factors particular to each implementation. Generally, this threshold value is selected in accordance with trial test runs of particular implementations to allow for "tuning" the threshold value. Using either heuristic, this received tag signal is then used (step 2016) in determining the tag distance.

Other embodiments of the signal processing hardware 1004 may include other hardware components in accordance with the particular digital signal processing requirements in a particular embodiment. Each particular embodiment may be implemented

in using a variety of combinations of hardware components, including, but not limited to, gate arrays and read-only-memory. Additionally, other embodiments of the signal processing hardware may be implemented as some combination of hardware and software in which the machine executable code may be executed on a computer system, such as the 5 microprocessor 1001 or the host computer 105 or yet another computer component included in the signal processing hardware 1004 as a dedicated processing unit. The components may vary with application and design choices associated with a particular implementation.

Another embodiment of the signal processing hardware includes hardware and/or software in addition to the correlator. This embodiment may perform, for example, all the 10 method steps of Figure 15B, rather than only some of the steps of Figure 15B, as in the embodiment with signal processing hardware that includes only the correlator.

The foregoing description sets forth a technique using an RLS method that affords a flexible and efficient way of filtering noise from a received signal as used in an article tracking system. Both multipath noise and tonal frequencies may be taken into 15 consideration as factors when using the foregoing techniques in various environments within which an article tracking system as described herein may operate.

The foregoing technique for filtering a received signal is scalable for use in applications with both large and small sample sets in a variety of different environments each having different "noise" considerations while generally providing a high degree of 20 accuracy in locating objects as required in article tracking systems as described herein.

Having described preferred embodiments of the invention, it will now become apparent to those of skill in the art that other embodiments incorporating its concepts may be provided. It is felt, therefore, that this invention should not be limited to the disclosed 25 embodiments, but rather should be limited only by the spirit and scope of the appended claims.

What is claimed is:

CLAIMS

1. A method of filtering received signals in a local tracking system to determine a received tag signal comprising:

- receiving a plurality of received signals;
- 5 determining a first range of the received signals;
- performing an autocorrelation function with phase shifts corresponding to the first range of received signals producing a corresponding magnitude for each received signal included in the first range;
- determining a second range of the received signals which is a subset of the first range, the second range approximating where in the first range the received tag signal is located;
- 10 performing recursive least-squares technique using the second range of received signals;
- producing a vector of filtered signals corresponding to the second range of received signals;
- 15 approximating a peak corresponding to a filtered portion of the received tag signal; and
- determining location of a tag associated with the received tag signal.

2. The method of Claim 1, wherein at least one of the received signals includes a noise component.

20 3. The method of Claim 1, wherein at least one of the received signals includes multipath noise.

4. The method of Claim 1, wherein at least one of the received signals includes a tonal frequency.

25 5. The method of Claim 4, wherein the tonal frequency is a jammer frequency generated by a device operating in the range of the article tracking system.

6. The method of Claim 1, wherein the article tracking system is used to track items within a range of a low-powered tag.

30 7. The method of Claim 1, wherein the local tracking system is used to track items indoors in an environment having multipath and tonal frequency noise.

8. The method of Claim 1, wherein the step of determining a second range of received signals includes:

using a rising edge detection technique to approximate the location of the received tag signal.

5 9. The method of Claim 1, wherein the step of determining a second range of received signals includes:

using a threshold detection technique to approximate the location of the received tag signal.

10 10. The method of Claim 1, wherein the step of determining a second range of received signals includes:

using a peak detection technique to approximate the location of the received tag signal.

11. The method of Claim 1, wherein the size of the second range is twelve.

12. The method of Claim 9, wherein the size of the second range is twelve, and step 15 of determining the second range includes:

determining a reference point approximating the received tag signal;

determining a starting point of the second range which is three samples times prior to the reference point; and

20 determining an ending point of the second range which is eight samples subsequent to the reference point.

13. The method of Claim 1, wherein said step of performing autocorrelation is performed in hardware using a field programmable gate array.

14. The method of Claim 1, wherein the step of approximating a peak 25 corresponding to a filtered portion of the received tag signal includes:

determining the largest magnitude of the received signals in the second range;

determining a threshold equal to 62.5% of the largest magnitude;

determining a first one of the signals in the second range which is first in time to equal or exceed the threshold;

30 determining one or more other received signals included in the second range; and

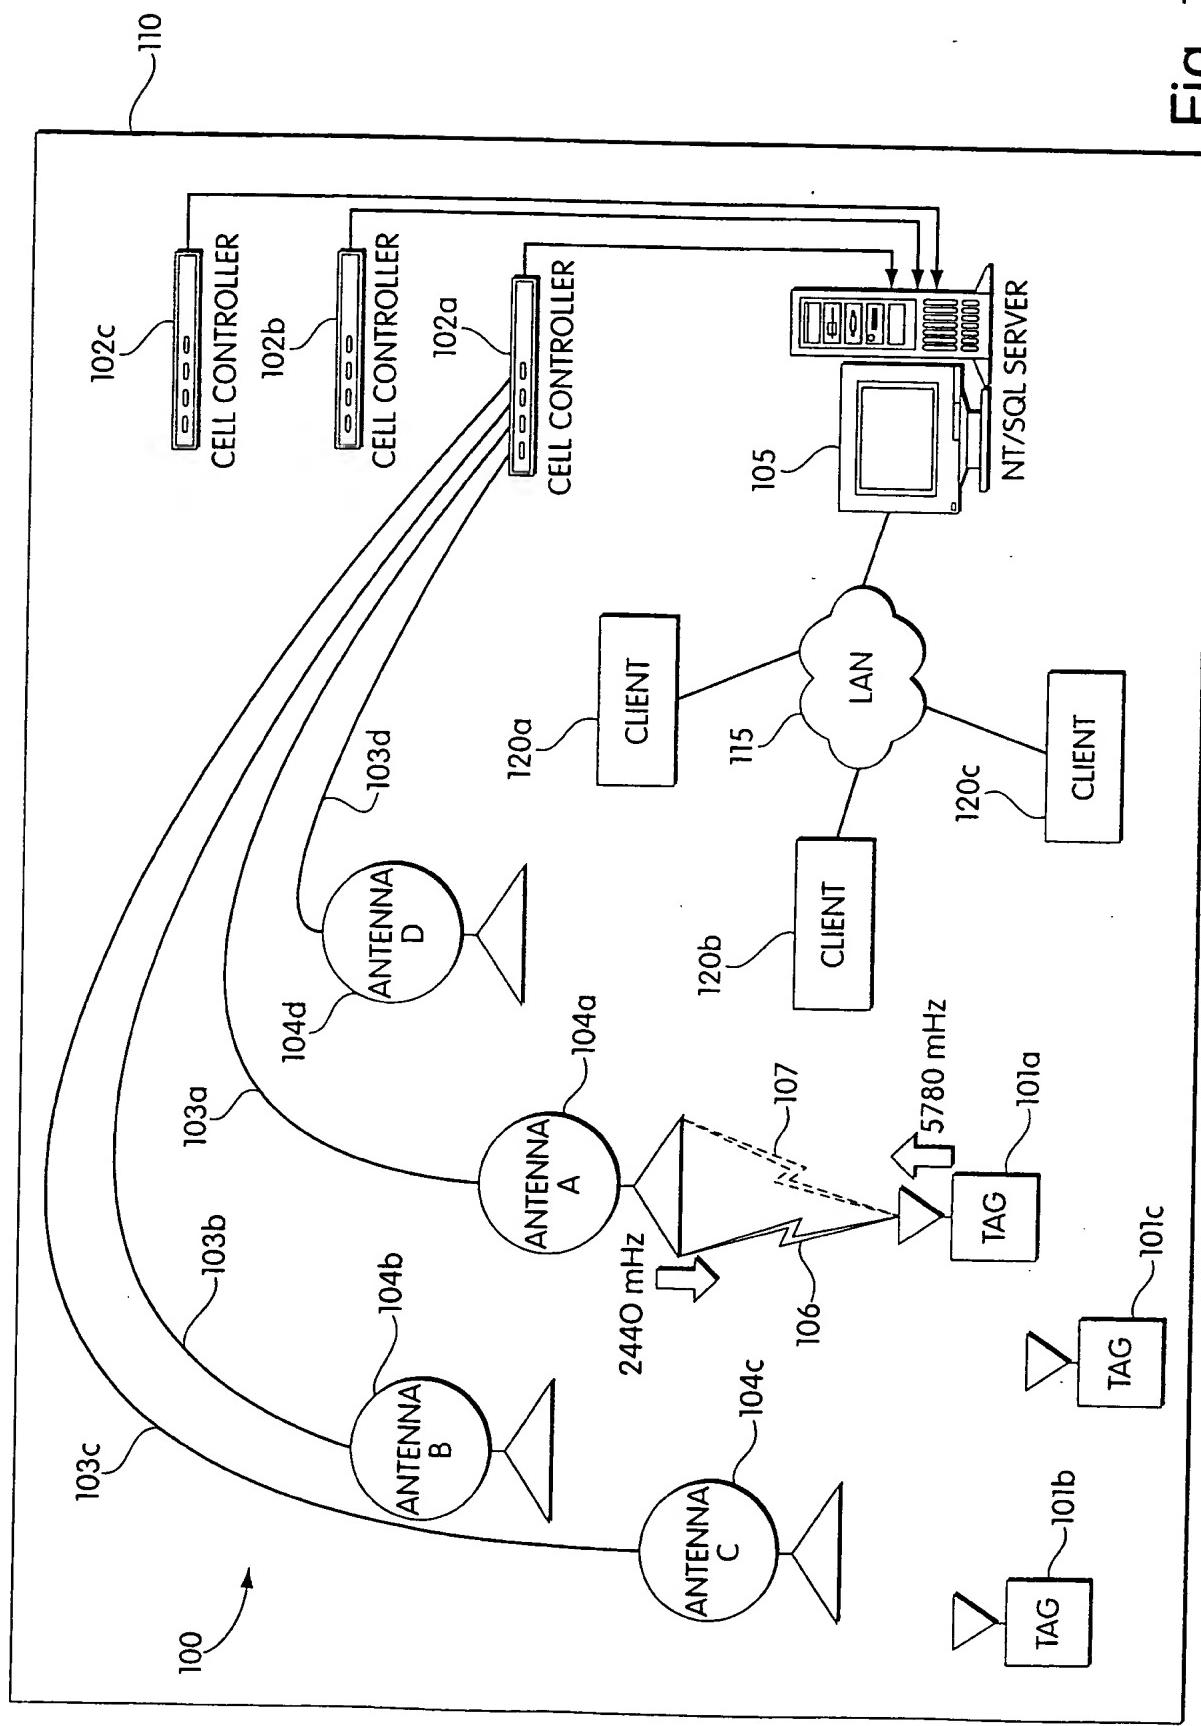
determining the received tag signal using said first and said one or more other received signals.

15. The method of Claim 14, wherein said one or more other received signals are subsequent to said first one of the received signals.

5 16. The method of Claim 14, wherein the step of determining the received tag signal using said first and said one or more other received signals includes:

using a weighted average of said first one and said one or more other received signals.

Fig. 1



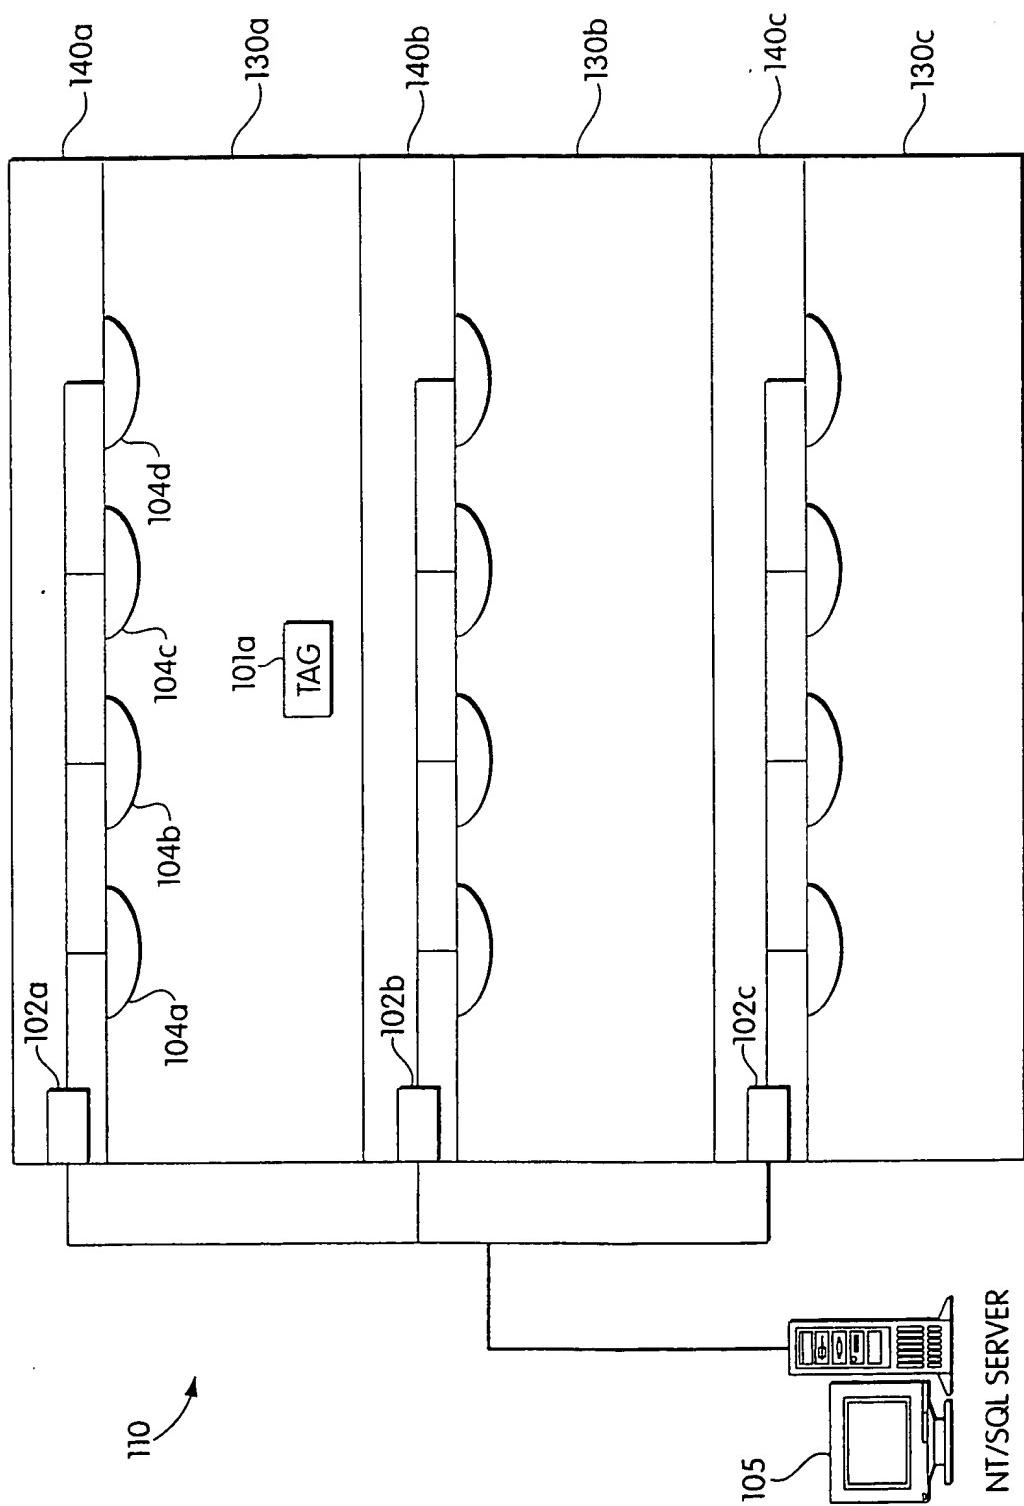


Fig. 2

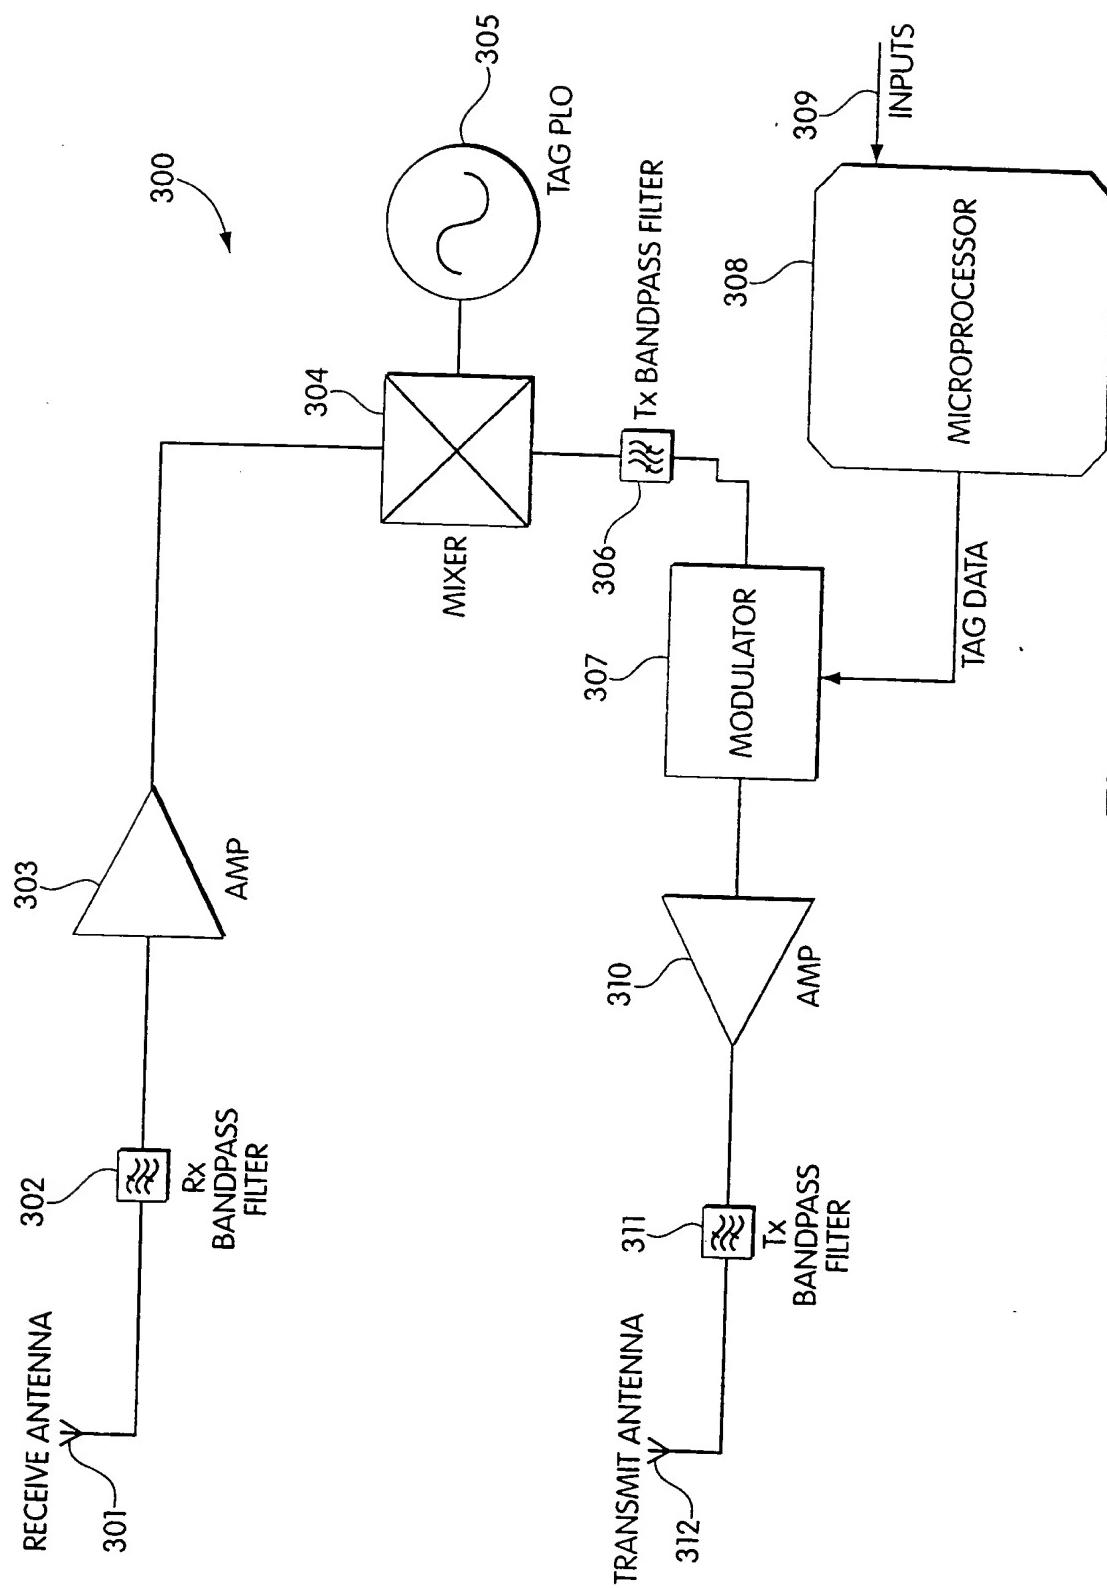


Fig. 3

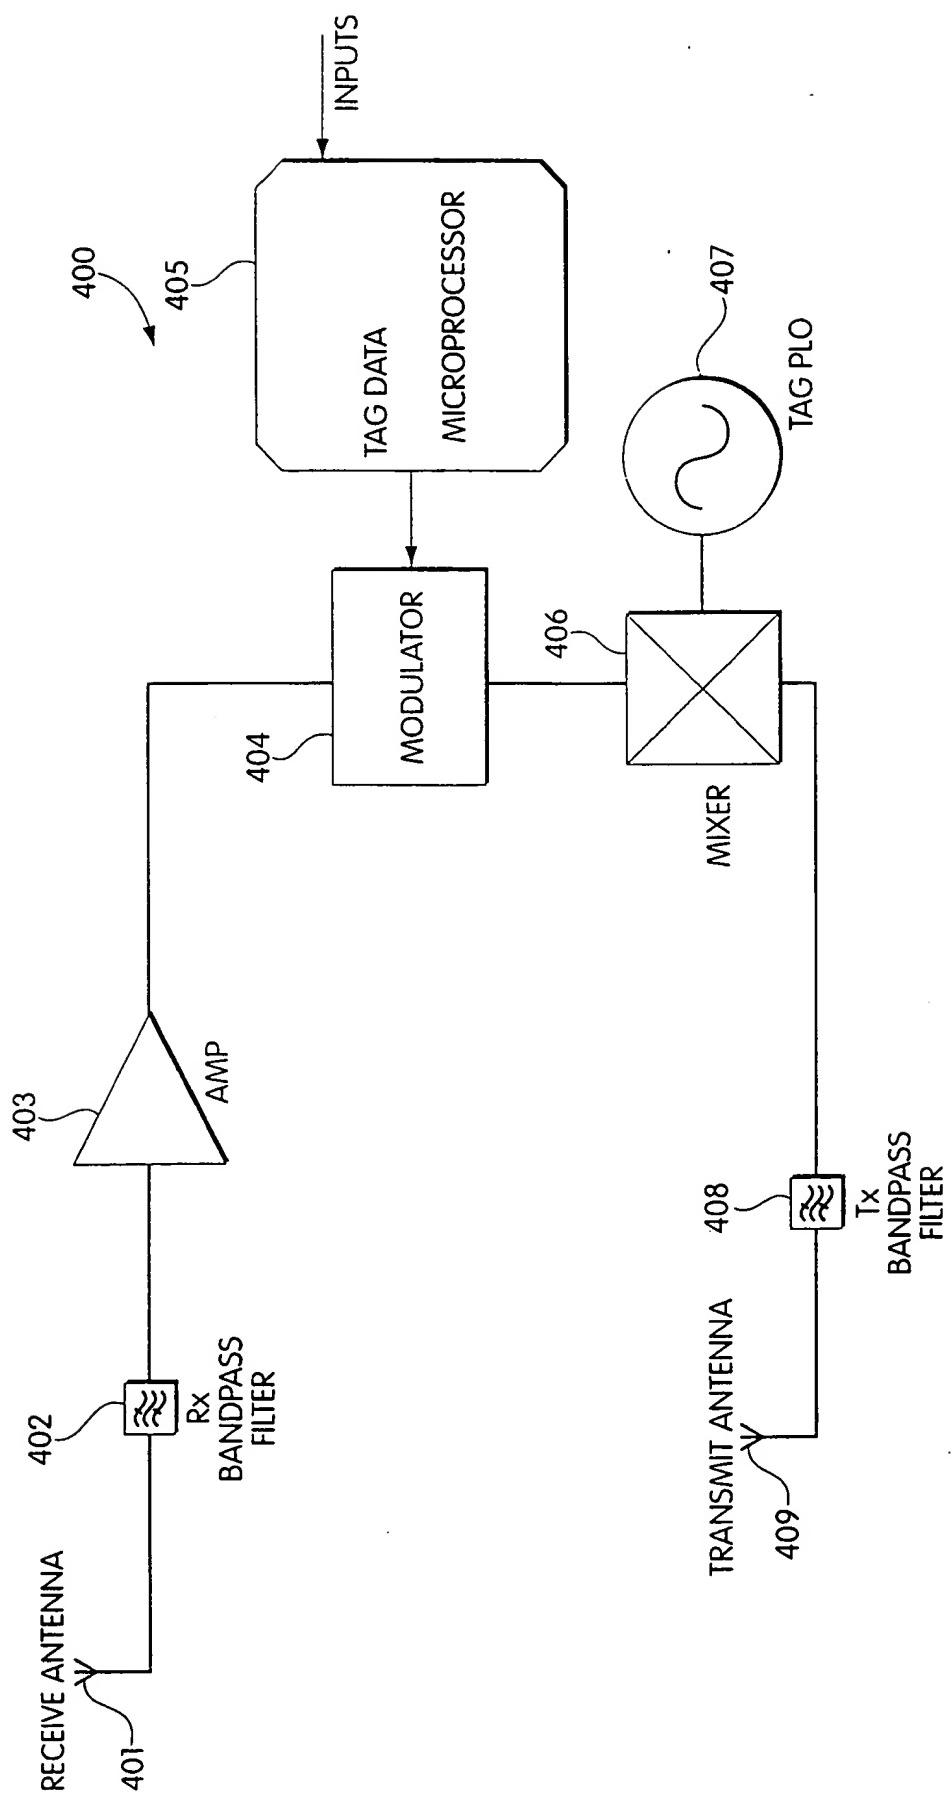


Fig. 4

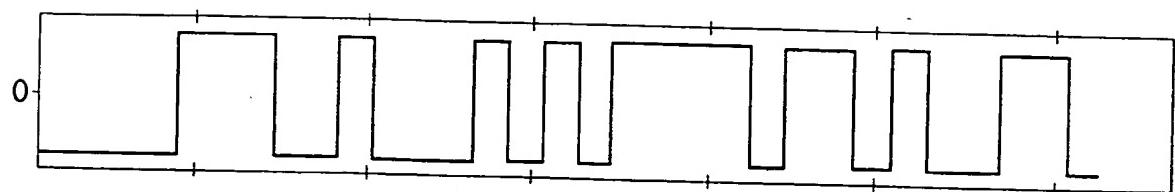


Fig. 5A

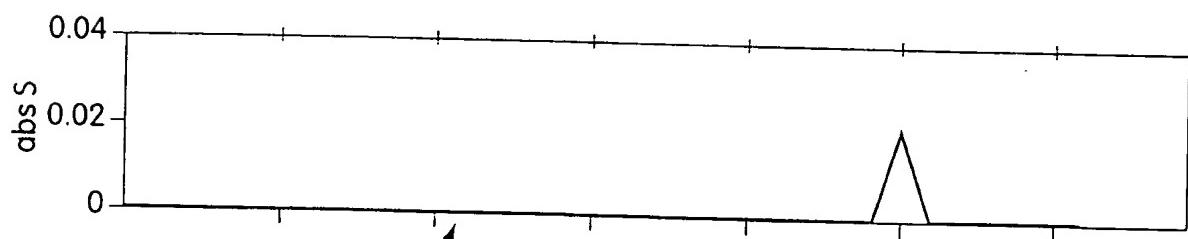


Fig. 5B

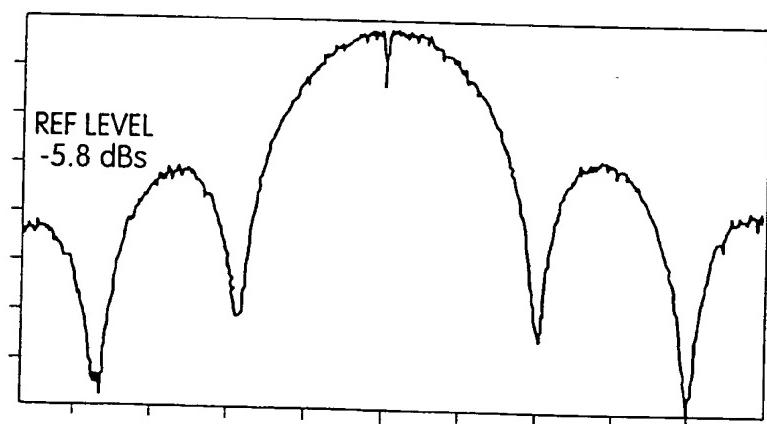


Fig. 5C

6/23

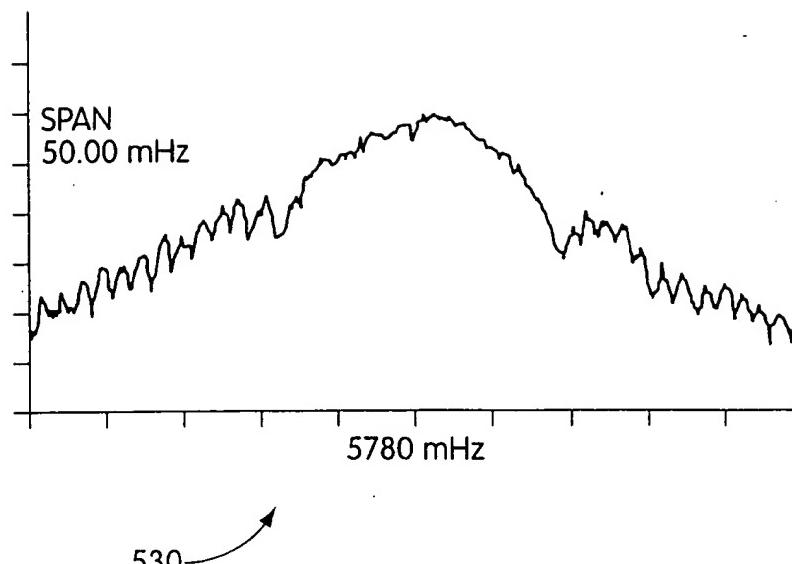


Fig. 5D

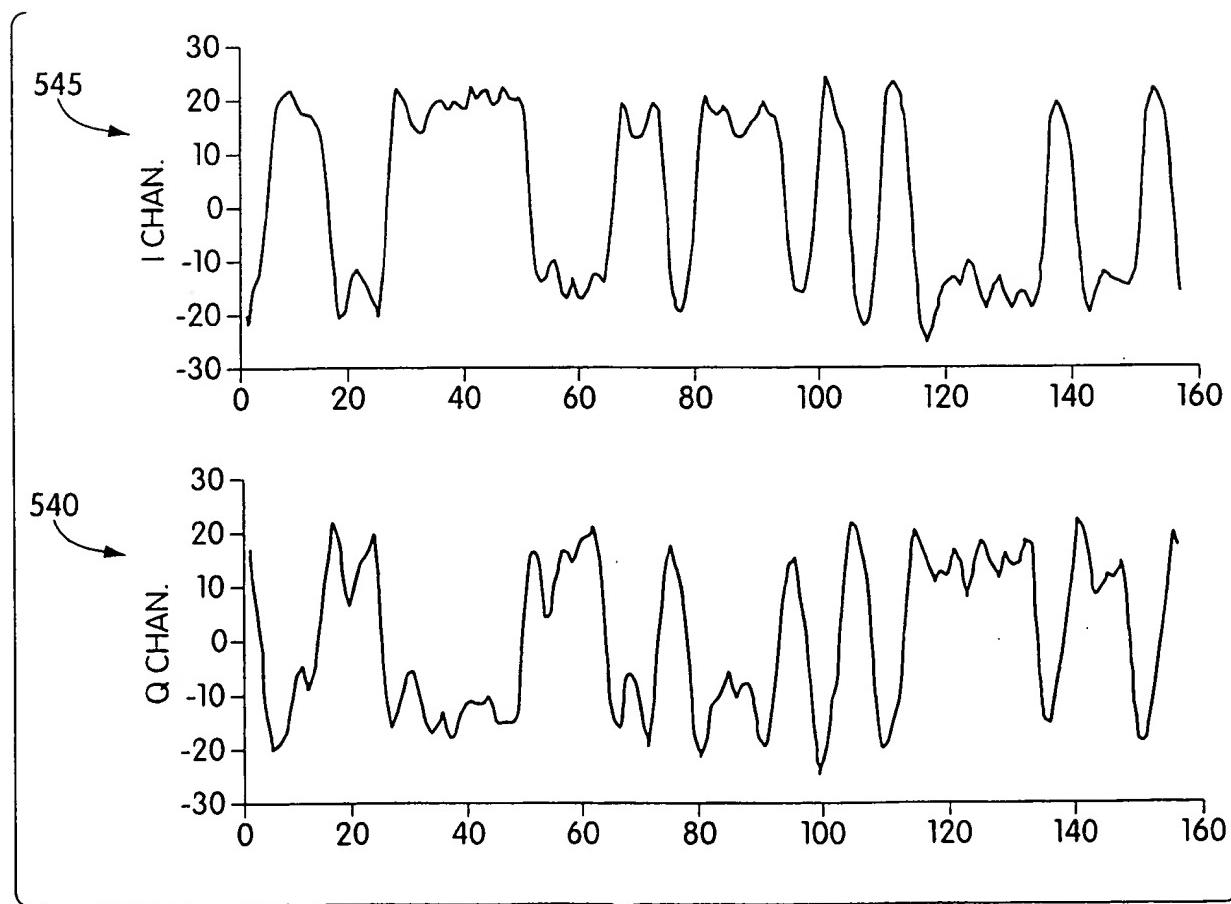


Fig. 5E

7/23

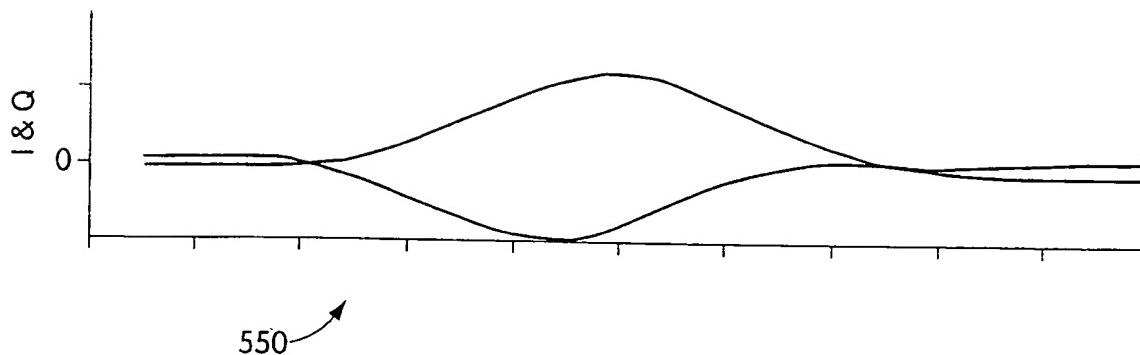


Fig. 5F

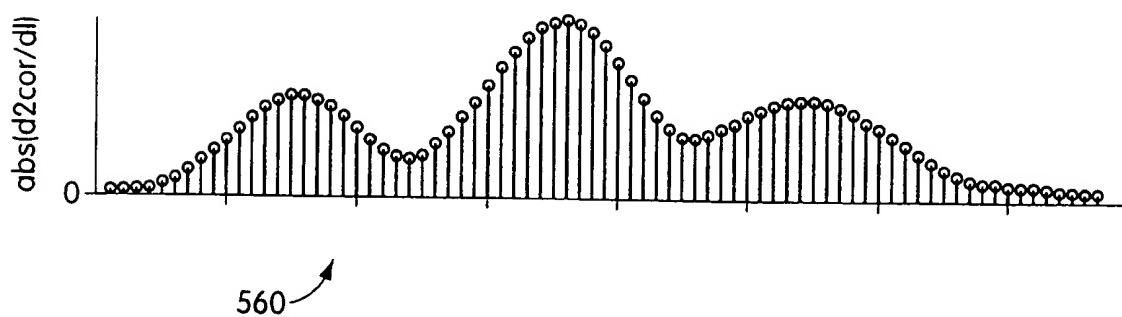
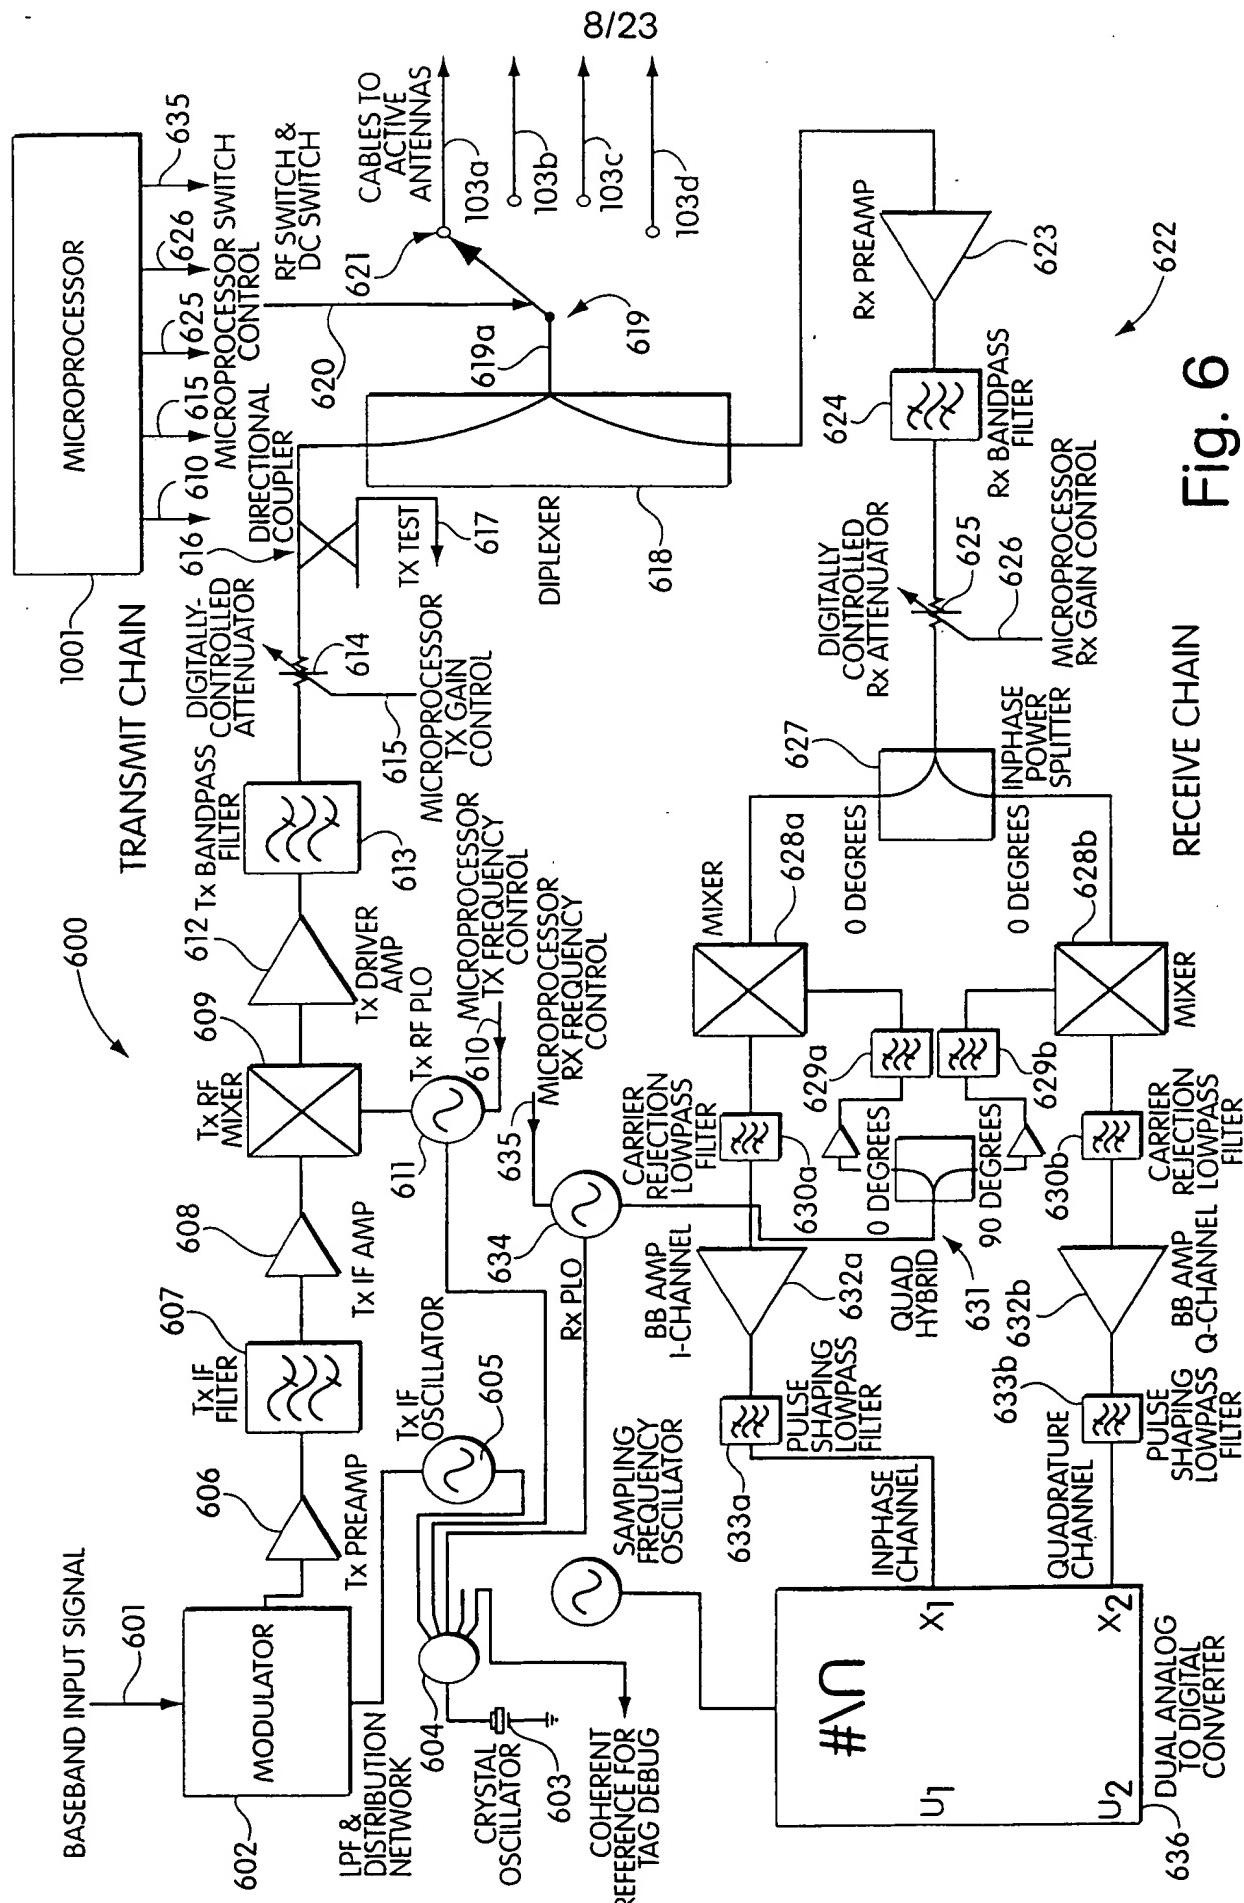


Fig. 5G



६
तं
प

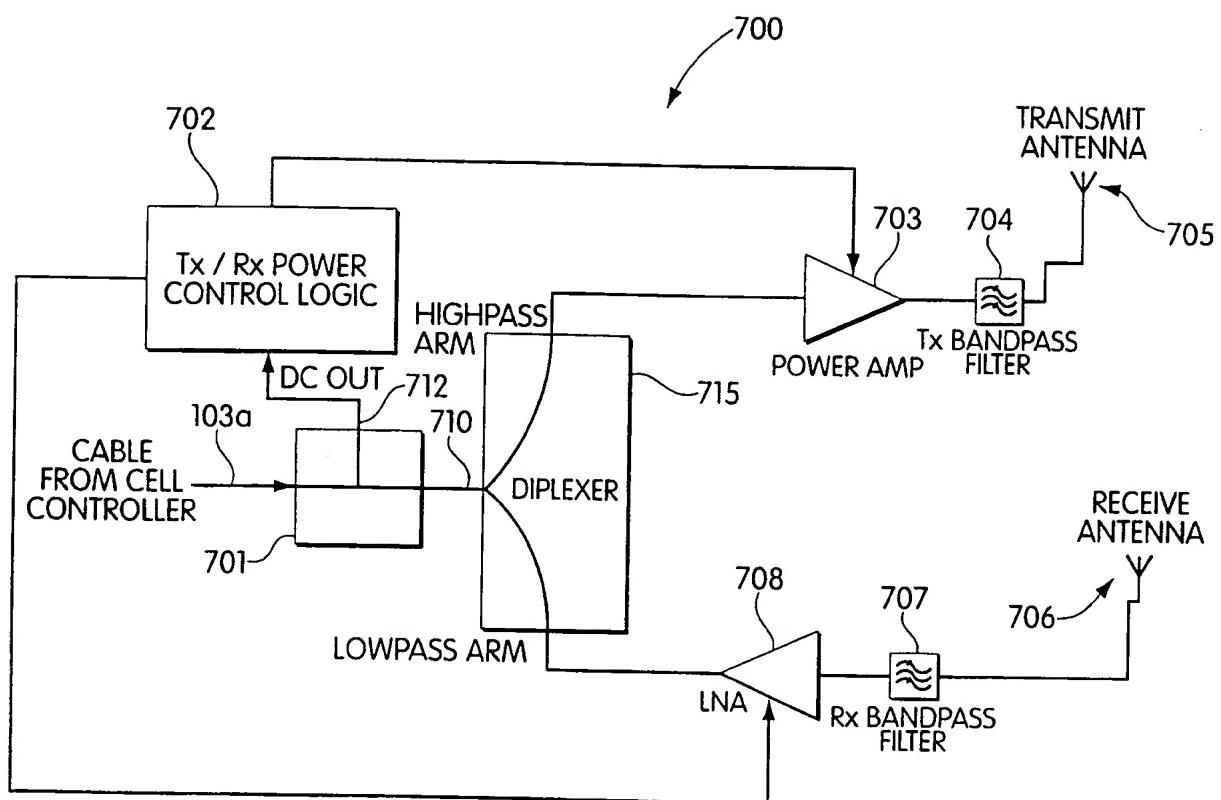


Fig. 7

10/23

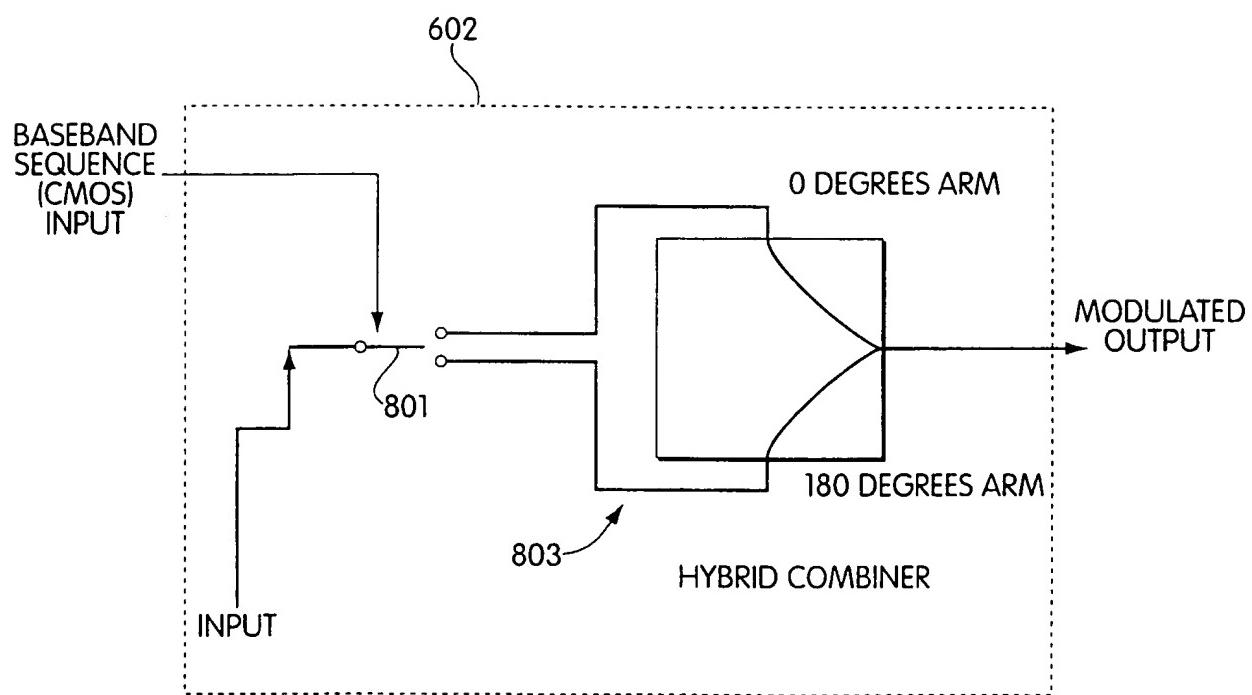


Fig. 8

11/23

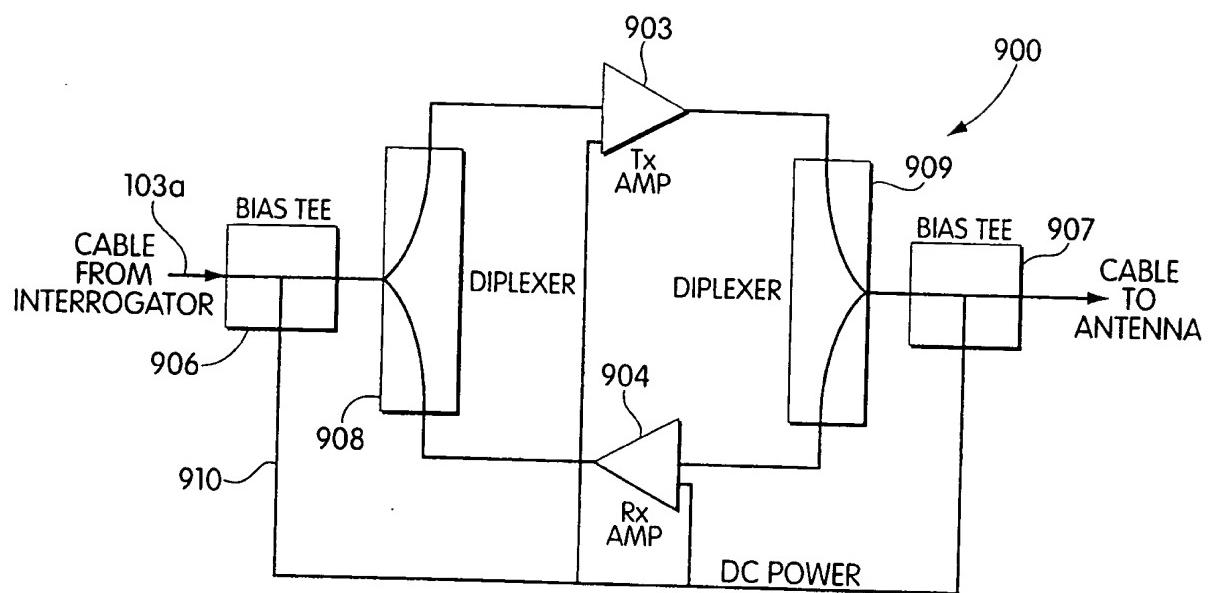


Fig. 9

12/23

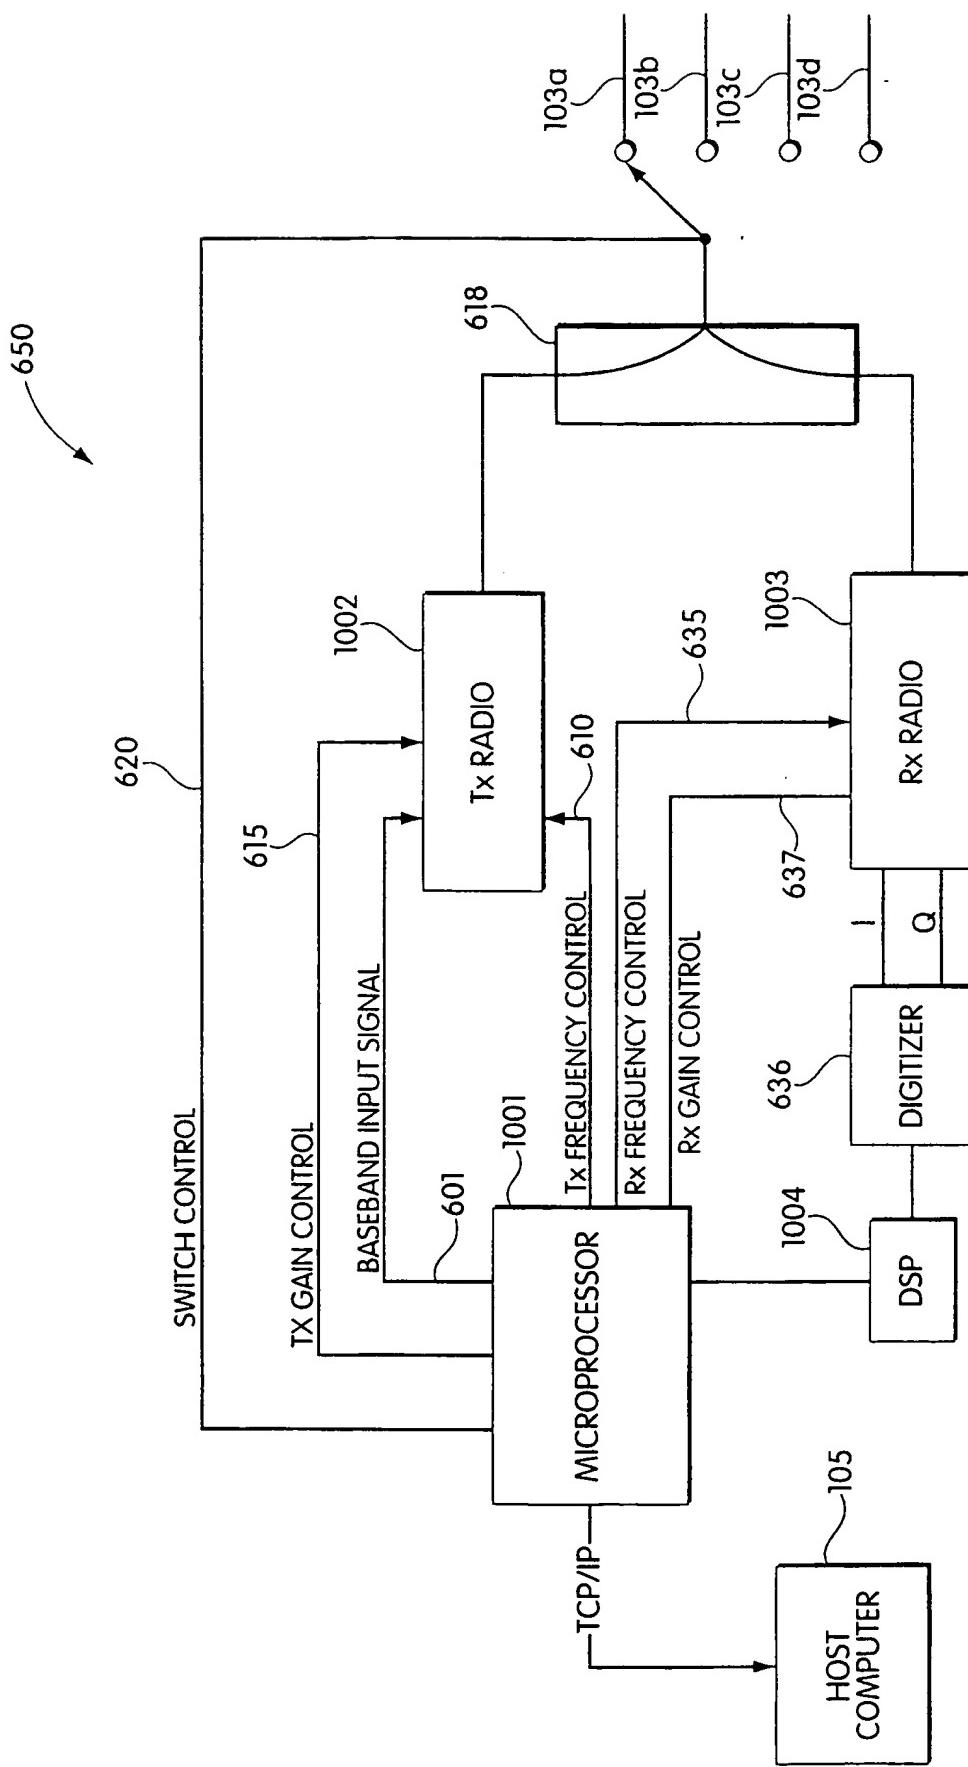


Fig. 10

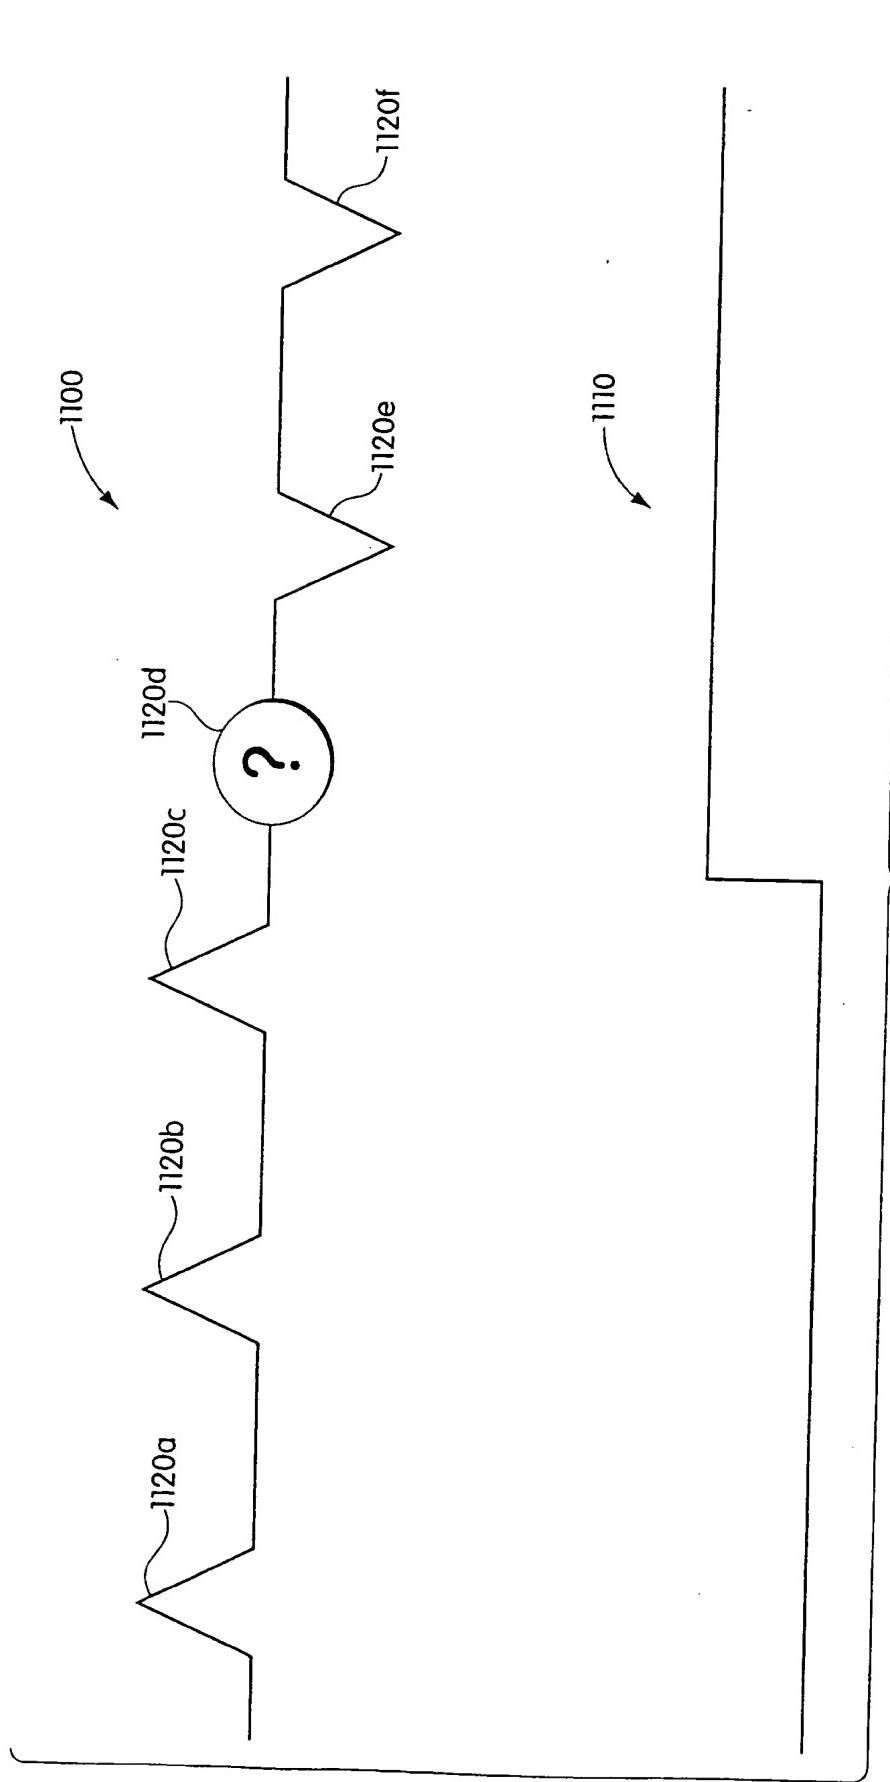


Fig. 11

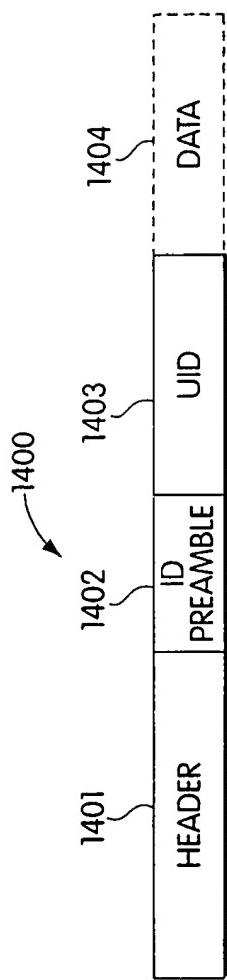


Fig. 12A

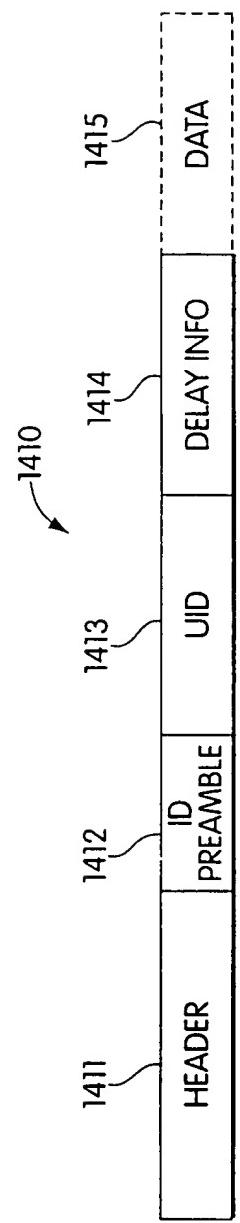


Fig. 12B

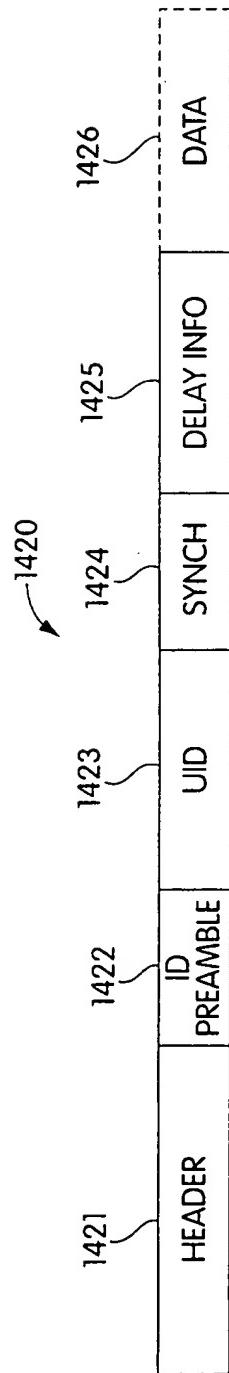


Fig. 12C

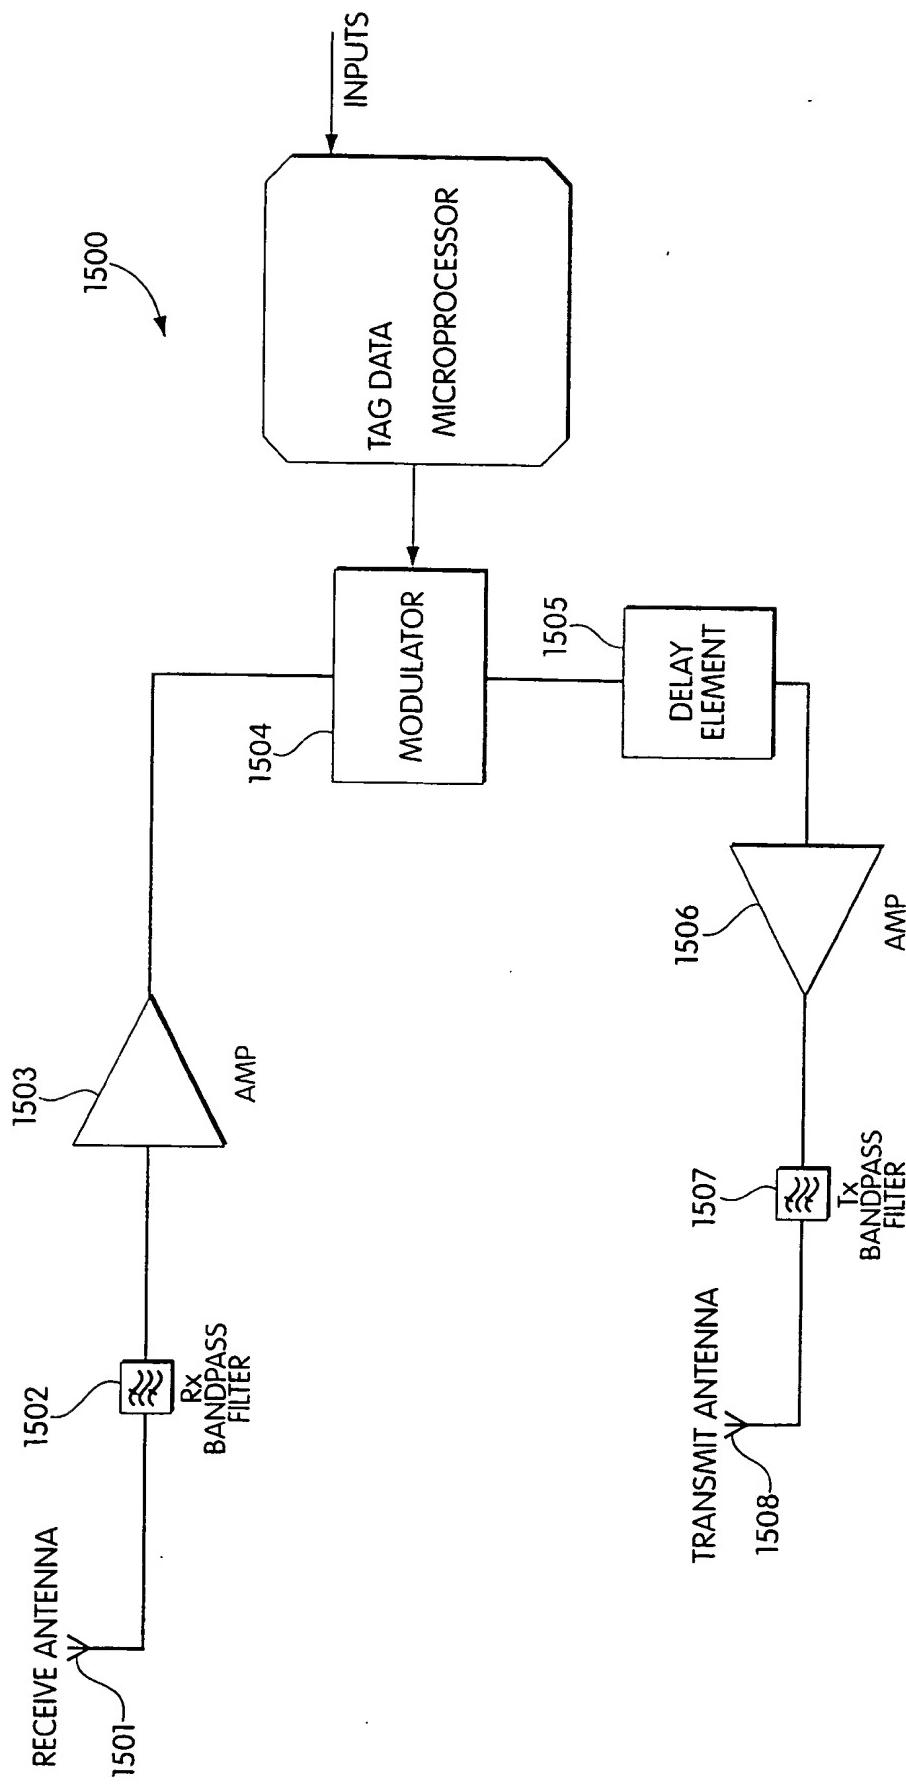


Fig. 13

16/23

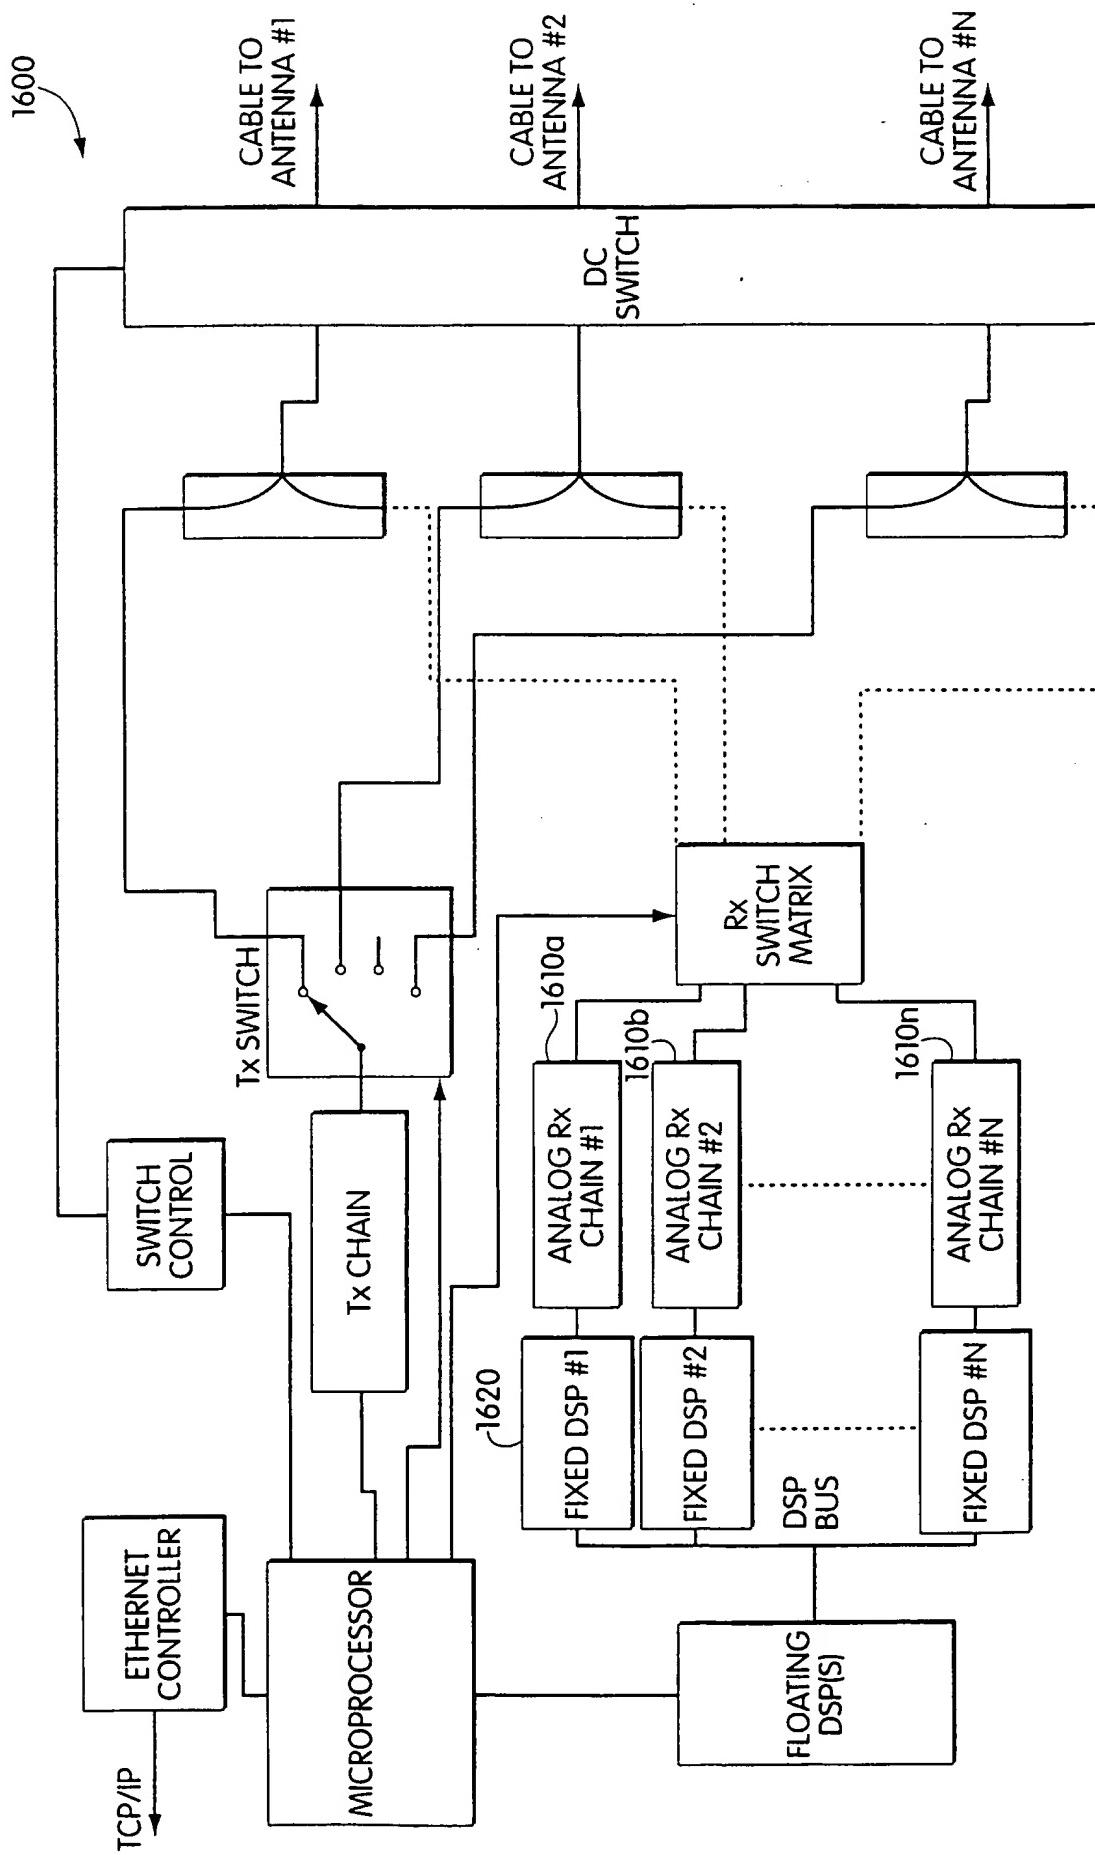


Fig. 14

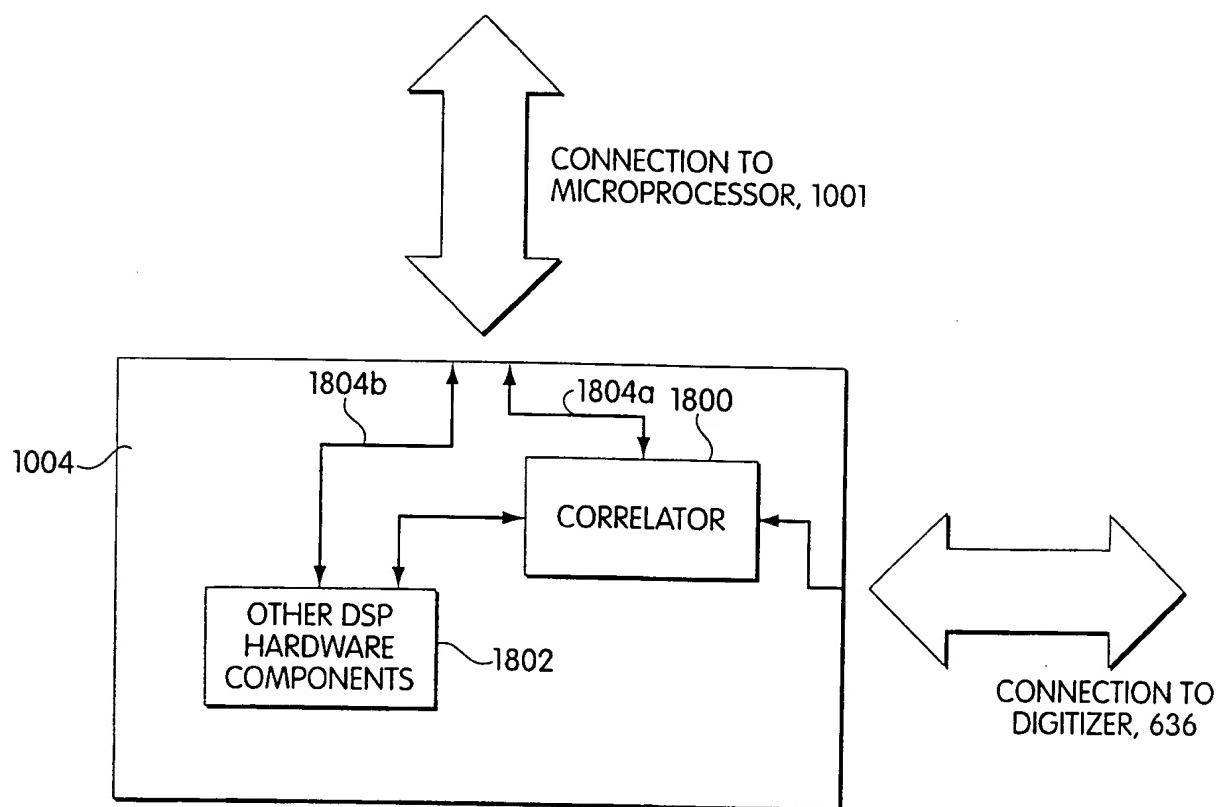


Fig. 15A

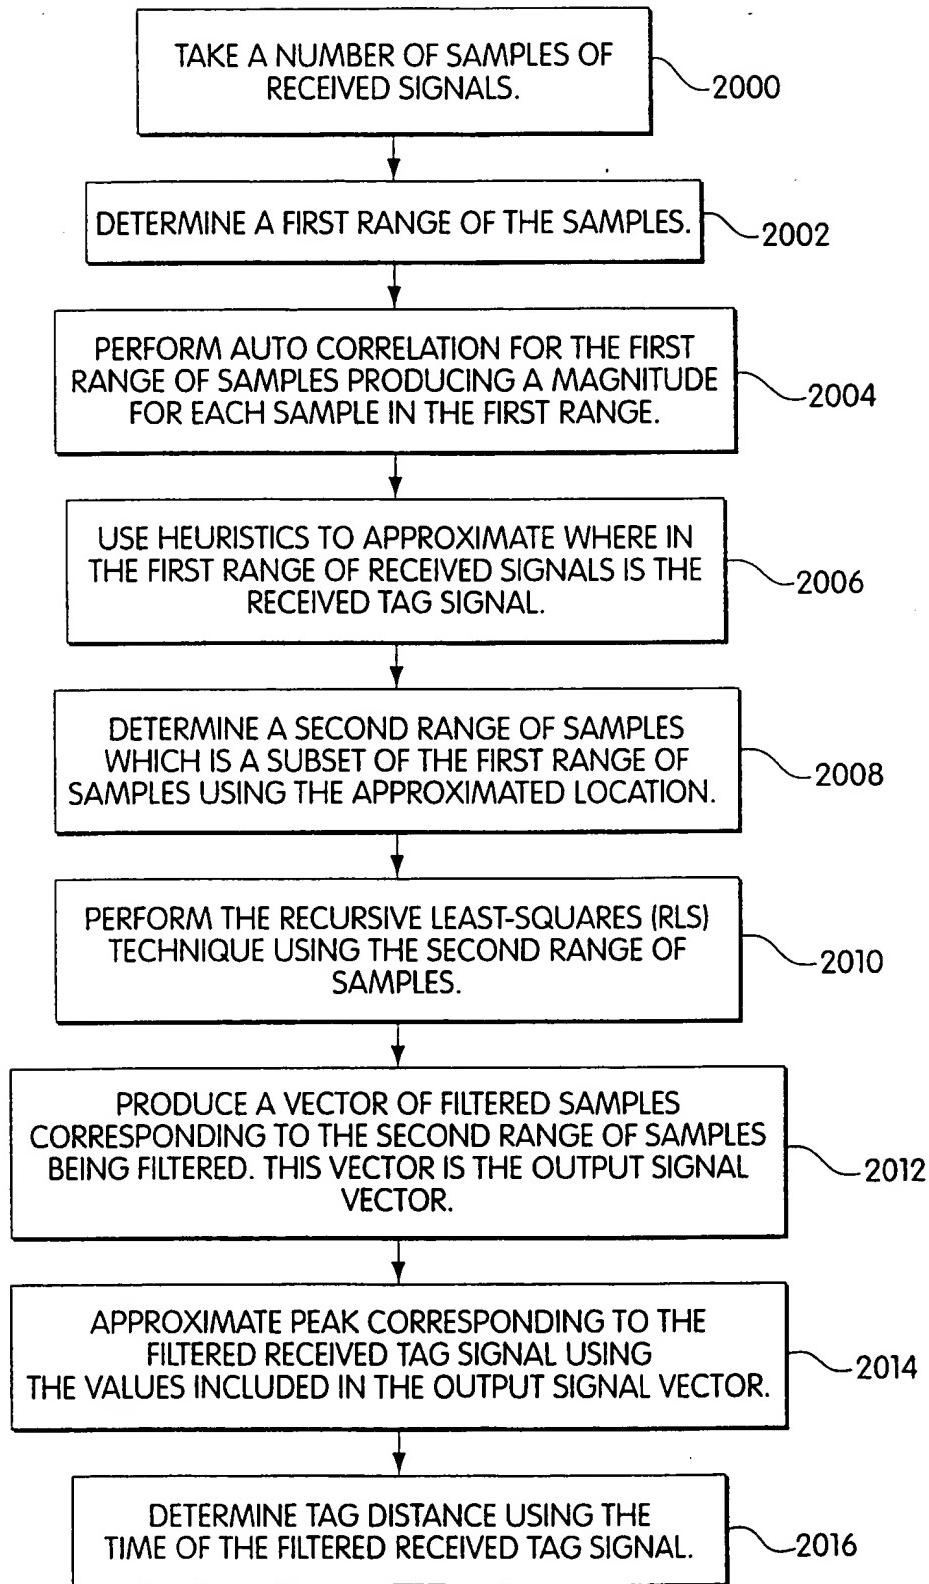


Fig. 15B

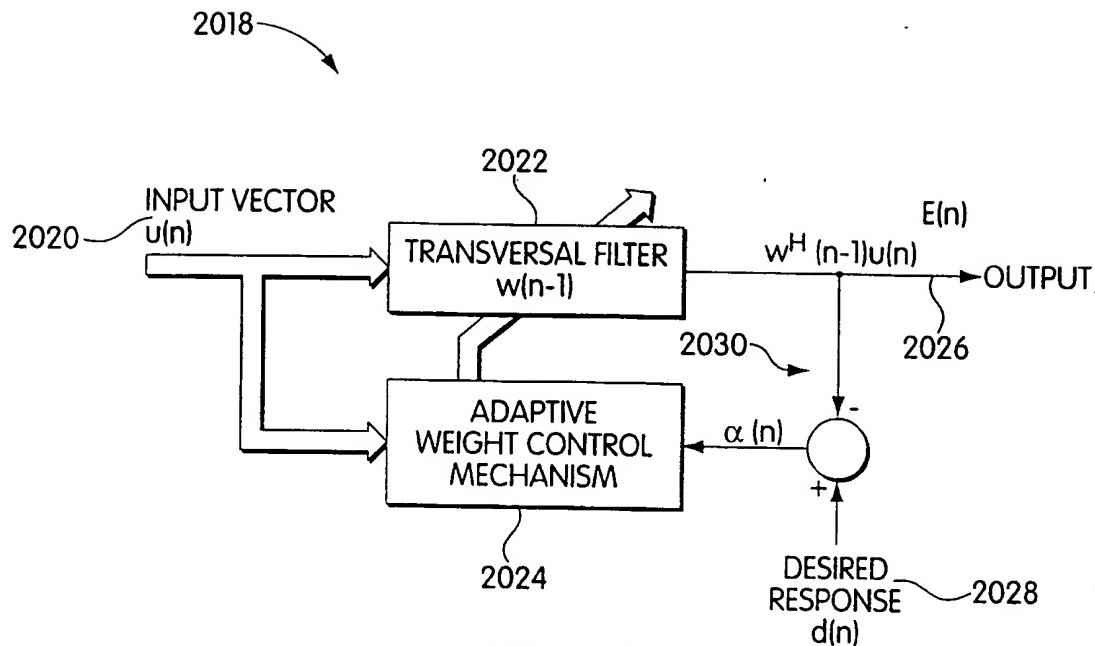


Fig. 16

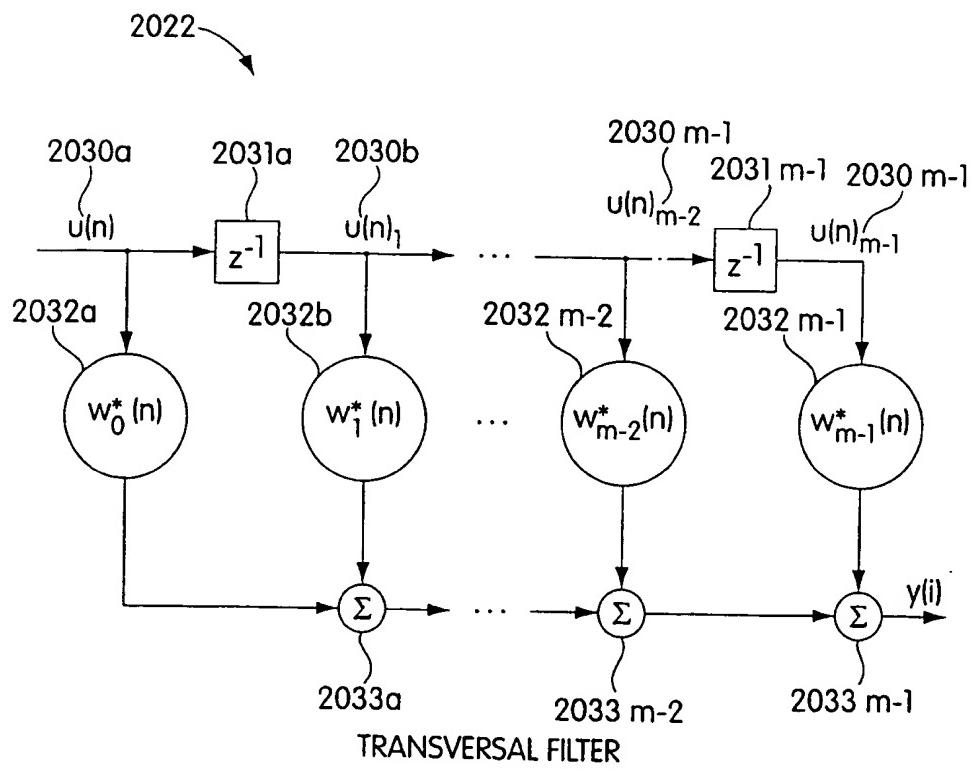


Fig. 17

20/23

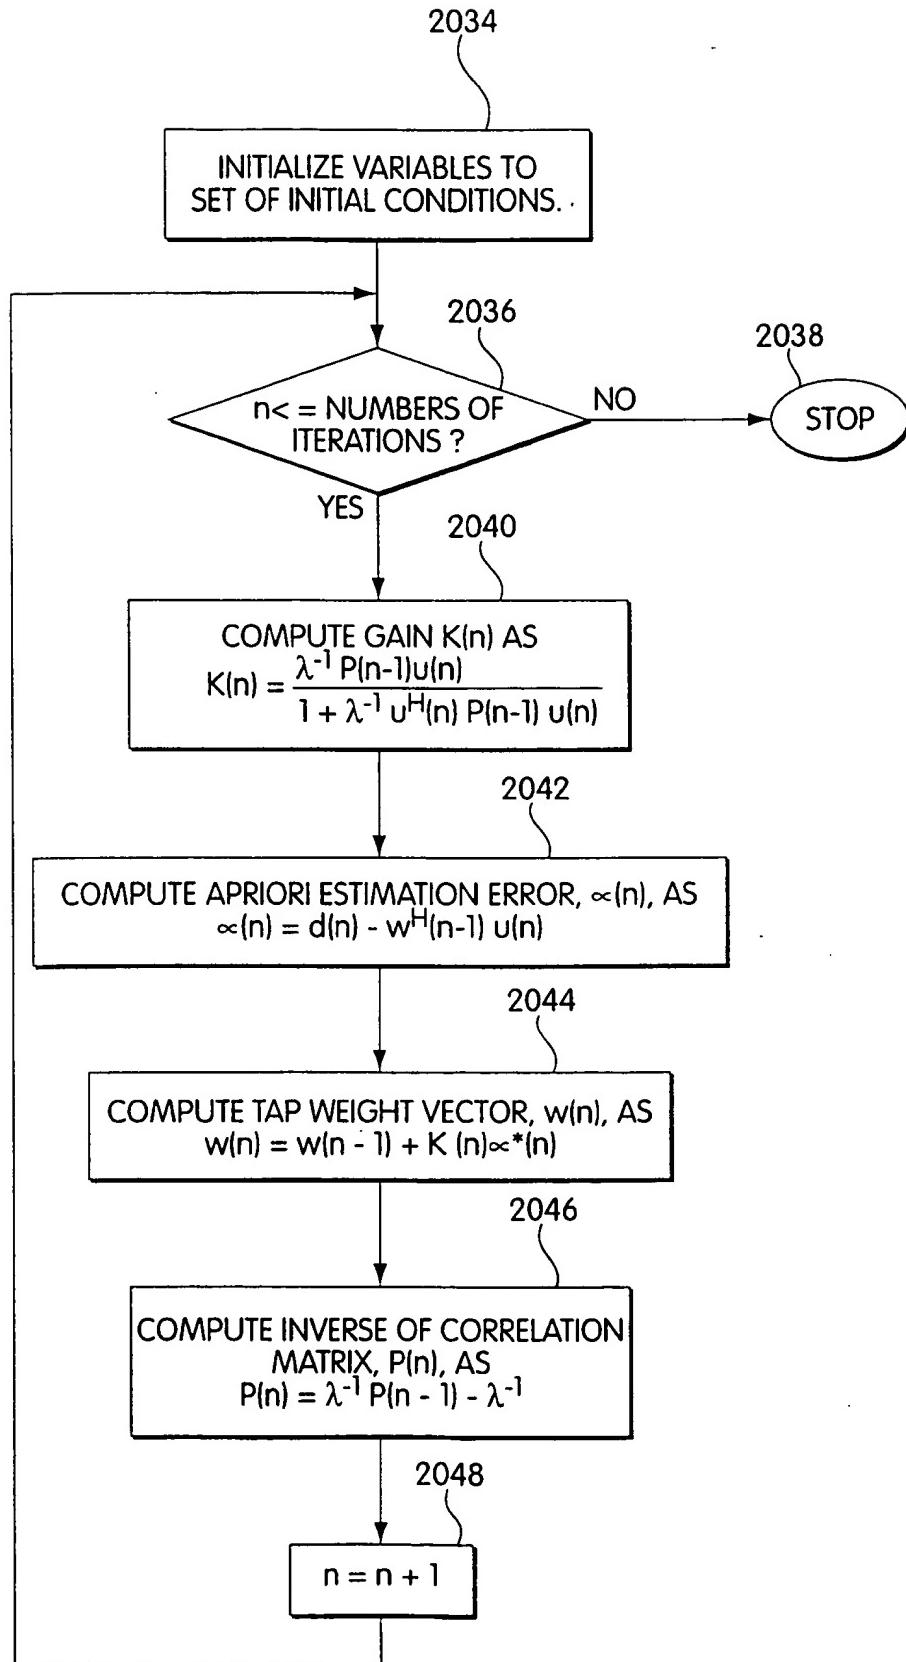


Fig. 18

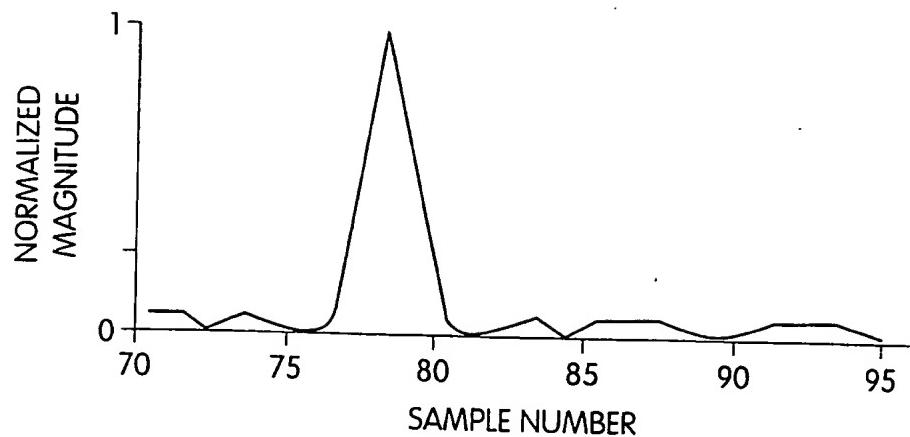


Fig. 19A

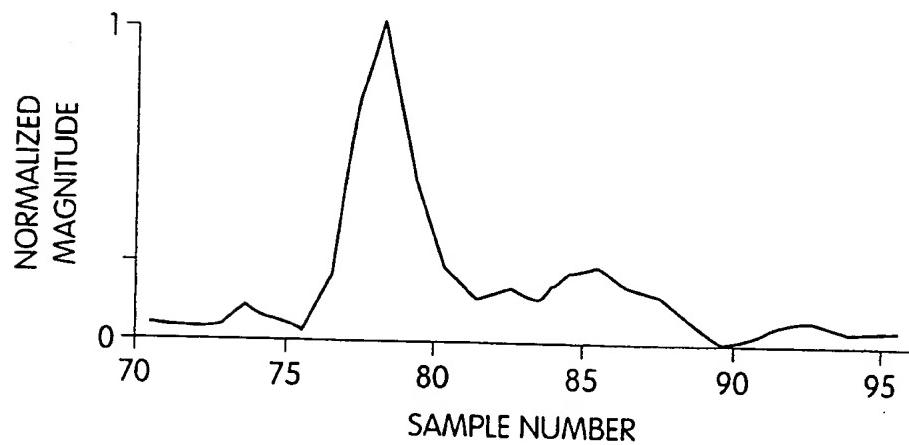


Fig. 19B

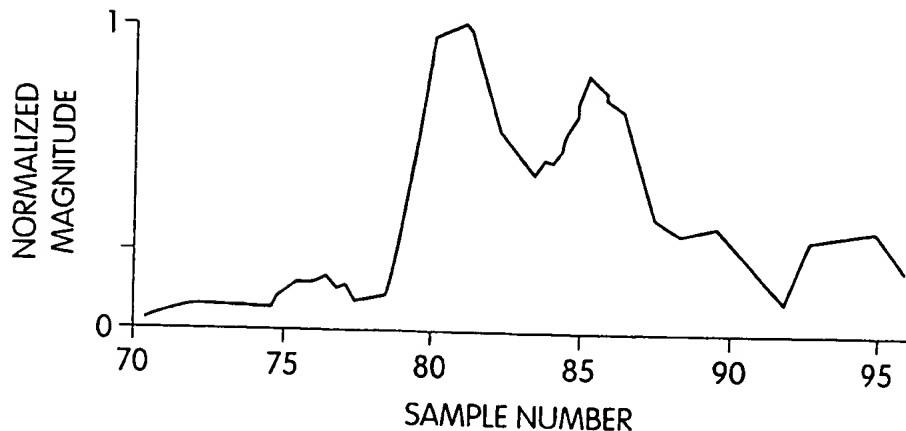


Fig. 19C

22/23

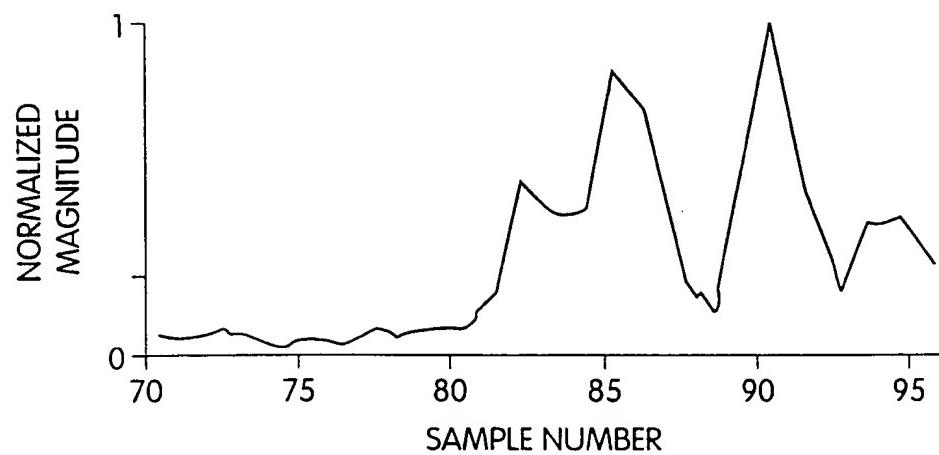


Fig. 19D

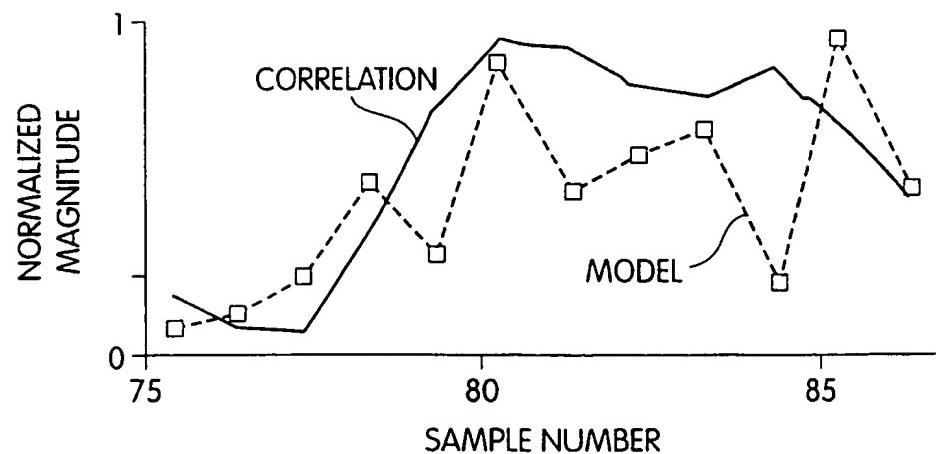


Fig. 19E

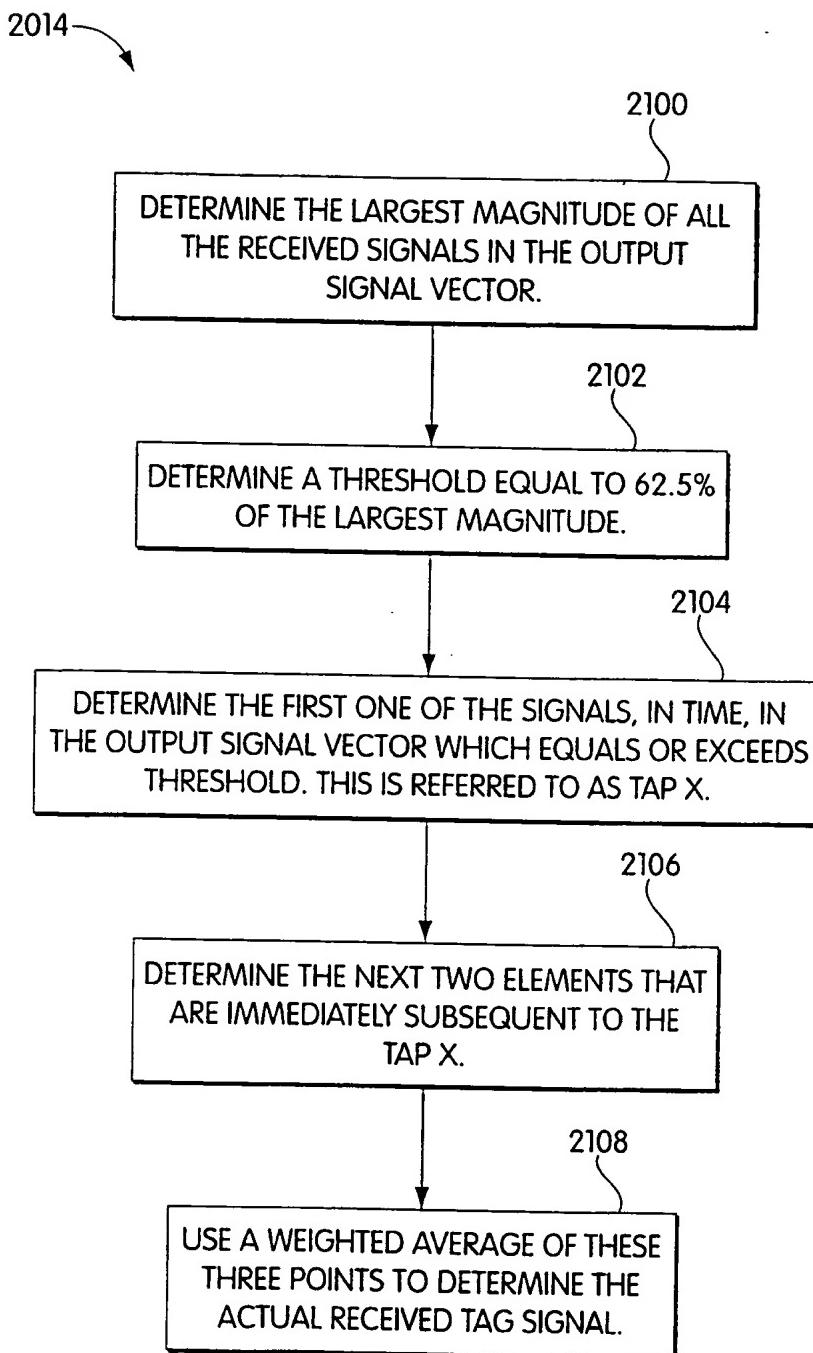


Fig. 20

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/02633

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :G08B 5/22; H04Q 5/22; H03H 7/30; G01S 13/74
US CL : 340/10.1, 10.2, 825.49; 375/233, 234; 342/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 340/10.1, 10.2, 825.49; 375/233, 234; 342/44.

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

West

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 4,818,998 A (APSELL et al) 04 April 1989, see entire document	1-16
Y	US 5,274,670 A (SERIZAWA et al) 28 December 1993, see entire document.	1-16
Y	US 5,285,480 A (CHENNAKESHU et al) 08 February 1994, see entire document.	1-16

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"B" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

22 APRIL 2000

Date of mailing of the international search report

16 MAY 2000

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Faxsimile No. (703) 305-3230

Authorized officer
[Signature]
YVES DALENCOURT

Telephone No. (703) 308 - 8547

This Page Blank (uspto)